

Queue-Dispatch Asynchronous Systems

Gilles Geeraerts Alexander Heußner
Jean-François Raskin

Université Libre de Bruxelles – Belgium

October 18, 2012

Abstract

To make the development of efficient multi-core applications easier, libraries, such as Grand Central Dispatch, have been proposed. When using such a library, the programmer writes so-called *blocks*, which are chunks of codes, and dispatches them, using *synchronous* or *asynchronous* calls, to several types of waiting queues. A scheduler is then responsible for dispatching those blocks on the available cores. Blocks can synchronize via a global memory. In this paper, we propose Queue-Dispatch Asynchronous Systems as a mathematical model that faithfully formalizes the synchronization mechanisms and the behavior of the scheduler in those systems. We study in detail their relationships to classical formalisms such as pushdown systems, Petri nets, fifo systems, and counter systems. Our main technical contributions are precise worst-case complexity results for the Parikh coverability problem and the termination question for several subclasses of our model. We give an outlook on extending our model towards verifying input-parametrized fork-join behaviour with the help of abstractions.

1 Introduction

The computing power delivered by computers has followed an exponential growing rate the last decades. One of the main reasons was the steady increase of the CPU clock rates. This growth, however, has come to an end a few years ago, because further increasing the clock rate would incur major engineering challenges related to power dissipations. In order to overcome this and meet the continuous need for more computing power, multi-core CPU's have been introduced and are now ubiquitous. However, in order to harness the power of multiple cores, software applications need to be fundamentally modified and the programmers now have to write programs with parallelism in mind. But writing parallel programs is a notoriously difficult and error prone task. Also, writing *efficient* and *portable* parallel code for multi-core platforms is difficult, as the number of available cores will vary greatly from one platform to another,

Parikh coverability		queue types		
		concurrent	serial	both
dispatch	synchron.	EXPTIME-C	PSPACE-C	EXPTIME-C
	asynchron.	EXPSpace-C	\nexists	(\nexists)
	both	\nexists	(\nexists)	(\nexists)

Table 1: QDAS Verification Problems (\nexists : “undecidable”, parentheses: directly derivable)

and might also depend on the current load, the energy management policy, and so forth.

In order to alleviate the task of the programmer, several high level programming interfaces have been proposed, and are now available on several operating systems. A popular example is *Grand Central Dispatch*, GCD for short, a technology that is present in MacOS X (since 10.6), iOS (since version 4), and FreeBSD. In GCD, the programmer writes so-called *blocks* which are chunks of codes, and send them to *queues*, together with several dependency constraints between those blocks (for instance, one block cannot start before the previous one in the queue has finished). The scheduler is then responsible for dispatching those blocks on the available cores, through a thread pool that the scheduler manages (thereby avoiding the explicit and costly creation/destruction of threads by the programmer that is in addition extremely error-prone).

So far, to the best of our knowledge, no formal model has been proposed for systems relying on GCD or similar technologies, making those programs *de facto* out of reach of current verification methods and tools. This is particularly unfortunate as the control structure of such programs is rich and may exhibit complex behaviors. Indeed, the state-space of such programs is infinite even when types of variables are abstracted to finite domains of values. This is not surprising as asynchronous calls and recursive synchronous calls can send an unbounded number of blocks to queues. Also, those programs are, as any parallel program, subject to concurrency bugs that are difficult to detect using testing only.

Contributions In this paper, we introduce *Queue-Dispatch Asynchronous Systems*, QDAS for short, as a formal model for programs written using libraries such as GCD. Our model is composed of *blocks*, that are finite transition systems with finite data-domain variables that can do *asynchronous* (non-blocking) and *synchronous* (blocking) calls to other blocks (possibly recursively). However, a call does not immediately trigger the execution of the callee: the block is inserted into a queue that can be either *concurrent* or *serial*. In concurrent queues, several blocks can be taken from the queue and executed in parallel, while in serial queues, a block can be dequeued only if the previous block in the queue has completed its execution. Queues are maintained with a fifo policy.

To formalize configurations of such systems, our formal semantics relies on *call task graph*, CTG for short, in which nodes model tasks that are either in queues or executing, and edges model dependencies between tasks and within queues.

We then study the decidability border for the *Parikh coverability problem* and the *termination problem* on several subclasses of QDAS. Our results are summarized in Table 1. The *Parikh image* of a CTG is an abstraction that counts for each type and state of blocks the number of occurrences in the CTG and the *Parikh coverability* problem asks for the reachability of a CTG that contains at least a given number of blocks of each type that are in a given set of states. Not surprisingly, this problem is undecidable for QDAS, but we identify several subclasses for which the problem is decidable. For those decidable cases, we characterize the exact complexity of the problem.

The main positive decidability results with precise complexity are as follows: First, we show that QDAS with *only* synchronous calls are essentially equivalent to pushdown systems with finite domain data-variables, and we show that the Parikh coverability problem is EXPTIME-C for synchronous concurrent QDAS (Theorem 1). Second, for synchronous QDAS with only serial queues, the problem is PSPACE-C (Theorem 2). Third, we show that QDAS with *only* asynchronous calls and *only* concurrent queues are essentially equivalent to lossy Petri nets and show that the Parikh coverability problem is EXPSPACE-C for that class (Theorem 3). This decidability border is precise as we show that if we allow either (i) asynchronous calls with synchronous queues, or (ii) synchronous and asynchronous calls with concurrent queues, then the Parikh coverability problem becomes undecidable (Theorem 4 and Theorem 5). The previous proof’s ideas allow to derive similar results for termination wrt. the subclasses of QDAS. The *termination* problem asks given a QDAS whether all its executions are finite.

We enhance up our results by presenting an extension of QDAS with an explicit fork/join construct that, in addition, is parametrized by the input. As Parikh coverability and termination lifted to this setting are undecidable, we propose two over-approximations that allow for solutions in practice.

Remark: Due to the lack of space, detailed formal proofs are deferred to the appendix.

Related Works The basic model checking result for asynchronous programs is the EXPSPACE-hardness for the control-state reachability problem obtained by making formal a link with *multi-set pushdown systems* (MPDS). The underlying two basic ideas are : (i) to untangle the call stack and the storage of pending asynchronous calls by imposing that the next call in a serialized execution-equivalent program is only processed when the call stack is empty; and (ii) to only count the number of pending calls for each block while the call stack is non-empty. The original reduction in [17] is based on Parikh’s theorem and derives the lower bound from a Petri net reachability problem [8]. A Parikh-less reduction was presented in [13] that relied on the convergence of an over- and under-approximation derived from interprocedural dataflow analysis.

The close relation between asynchronous programs and Petri nets can also be used to prove additional decidability results for liveness questions [11, 10]. The following results are based on a (polynomial-time) reduction of asynchronous systems to an “equivalent” Petri net or extension thereof: *fair* termination (i.e., testing whether each dispatched call terminates) is complete in EXPSpace, the boundedness question is decidable in EXPSpace (i.e., asking whether we can bound the number of pending calls), fair non-starvation (i.e., asking, when assuming fairness on runs, whether every pending call is eventually dispatched) is decidable. The authors also consider extensions of asynchronous programs with cancellation (i.e., an additional operation removing all pending instances of a block) and testing whether there is *no* pending instance of a given block. In the first case, they show reduction to the model to Petri nets with transfer arcs or reset arcs, in the second case they show reduction to Petri nets with one inhibitor arc. Multi-set pushdown automata are subsumed by *well-structured transition systems with auxiliary storage* and inherit their decidability results presented in [6, 7]. Analogously, one can show that termination, control-state maintainability, and simulation with respect to finite state systems are decidable for asynchronous programs.

All the models considered in the aforementioned publications do not consider causality constraints on the sequence of asynchronous dispatch calls, as would be necessary to model the fifo policies of GCD. However, this is possible with QDAS. A more detailed look on the differences between the model of [10] and the (fifo-less) subclass of asynchronous serial QDAS is presented in Section 4.

A series of parallel programming libraries and techniques is formalized in [3] with the help of *recursively parallel programs*. These allow to model fork/join based parallel computations based on a reduction to recursive vector addition systems with states. With respect to QDAS and asynchronous programming, recursively parallel programs only cover the classical asynchronous models presented above and not the advanced scheduling strategies for different queues that introduce more sophisticated behaviours.

2 Preliminaries

Grand Central Dispatch (GCD) is a technology developed by Apple [1, 2] that is publicly available at <http://libdispatch.macosforge.org/> under a free license. GCD is the main inspiration for the formal model of queue-dispatch asynchronous systems. In the following, we often present our examples as pseudo code using a syntax inspired by GCD. In the GCD framework, the programmer has to organize his code into *blocks*. During the execution of a GCD program, one or several *tasks* run in parallel, each executing a given block (initially, only the `main` block is running). Tasks can call (or *dispatch* in the GCD vocabulary) other blocks, either *synchronously* (the call is blocking), or *asynchronously* (the call is not blocking). A *dispatch* consists in inserting the block into a fifo *queue*. In our examples, we use the keywords `dispatcha` and `dispatchs` to refer to asynchronous and synchronous dispatches respectively. At any time, the sched-

```

1 global int const l,m,n
2 global int[l][m] matrix1, int[m][n] matrix2, int[l][n] matrix
3 global c_queue workqueue, s_queue semaphore, int count
4 block increase():
5     count = count + 1
6 block one_cell(int i, int j):
7     for k in range(m):
8         matrix[i][j] += matrix1[i][k] * matrix2[k][j]
9     dispatch_s(semaphore, increase())
10 def main():
11     // read input matrix1, matrix2
12     count = 0
13     for i in range(l):
14         for j in range(n):
15             dispatch_a(workqueue, one_cell(i,j))
16     wait(count = l*n)
17     // print the result

```

Figure 1: GCD (-like) program for parallel matrix multiplication

uler can decide to *dequeue* blocks from the queues and to assign them to tasks for execution. All queues ensure that the blocks are dequeued in fifo order, however the actual scheduling policy depends on the type of queue. GCD supports two types of queues: *concurrent queues* allow several tasks from the same queue to run in parallel, whereas *serial queues* guarantee that *at most one* task from this queue is running. In our examples, concurrent (or serial) queues are declared as global variables of type `c_queue` (`s_queue`). In addition, all blocks have access to the same set of *global variables* (in this work, we assume that the variables range over finite domains).

Example 1 *Let us consider the pseudo code in Fig. 1 that computes the product of two integer matrices `matrix1` and `matrix2` of constant size (l, m, n) in a matrix `matrix`. The `main` task forks a series of `one_cell` blocks. Each `one_cell` computes the value of a single cell of the result. The parallelism is achieved via the GCD scheduler, thanks to asynchronous dispatches on the concurrent queue `workqueue`. Asynchronous dispatches are needed to make sure that `main` is not blocked after each dispatch, and a concurrent queue allow all the `one_cell` block to run in parallel. The variable `count` is incremented each time the computation of a cell is finished and acts as a semaphore for the `main` block, to ensure that `matrix` contains the final result. As only reading and writing to a variable are atomic, we need to guarantee exclusive access of two consecutive operations on `count` (line 5). This is achieved by a dedicated block `increase` that is dispatched to the serial queue `semaphore`. As only `increase` blocks can increase `count`, this queue implicitly locks the access to the variable. Moreover, the synchronous dispatch in line 9 guarantees that a block terminates only after it has increased `count`.*

Basic Notations: Given a set S , let $|S|$ denote its cardinality. For an I -indexed family of sets $(S_i)_{i \in I}$, we write elements of $\prod_{i \in I} S_i$ in bold face, i.e.,

$\vec{s} \in \prod_{i \in I} S_i$. The i -component of \vec{s} is written $s_i \in S_i$, and we identify \vec{s} with the indexed family of elements $(s_i)_{i \in I}$. We use \cup to denote the disjoint union of sets. An *alphabet* Σ is a finite set of *letters*. We write Σ^* for the set of all *finite words*, over Σ and denote the empty word by ε . The concatenation of two words w, w' is represented by $w \cdot w'$. For a letter $\sigma \in \Sigma$ and a word $w \in \Sigma^*$, let $|w|_\sigma$ be the number of occurrences of σ in w . We use standard complexity classes, e.g., polynomial time (PTIME) or deterministic exponential time (EXPTIME), and mark completeness by appending “-C” (PSPACE-C).

Let \mathbb{D} be a finite **data domain** with an *initial element* $d_0 \in \mathbb{D}$, and let \mathcal{X} be a finite set of variables ranging over \mathbb{D} . A *valuation* of the variables in \mathcal{X} is a function $\mathbf{d} : \mathcal{X} \rightarrow \mathbb{D}$. An *atom* is an expression of the form $x = d$ or $x \neq d$, where $x \in \mathcal{X}$ and $d \in \mathbb{D}$. A *guard* is a finite conjunction of atoms. An assignment is an expression of the form $x \leftarrow v$, where $x \in \mathcal{X}$ and $v \in \mathbb{D}$. Let $\mathbf{guards}(\mathcal{X})$, $\mathbf{assign}(\mathcal{X})$ and $\mathbf{vals}(\mathcal{X})$ denote respectively the sets of all guards, assignments and valuations over variables from \mathcal{X} . Guards, atoms and valuations have their usual semantics: for all valuations \vec{d} of \mathcal{X} and all $g \in \mathbf{guards}(\mathcal{X})$, we write $\vec{d} \models g$ iff \vec{d} satisfies g .

A **pushdown system with data** is a pushdown system (see [4] for details) equipped with a finite set of variables \mathcal{X} over a finite domain \mathbb{D} . A configuration of a PDS with data is a pair (s, w, \vec{d}) where s is a control state, w is the stack content, and \vec{d} is a valuation of the variables

Proposition 1 *The reachability problem is EXPTIME-C for PDS with data.*

A **Petri net** (PN) is a tuple $N = \langle P, T, m_0 \rangle$ where P is a finite set of places, a *marking* of the places is function $m : P \rightarrow \mathbb{N}$ that associates, to each place $p \in P$ a number $m(p)$ of tokens, T is finite set of transitions, each transition $t \in T$ is a pair (I_t, O_t) where $I_t : P \rightarrow \{0, 1\}$ and $O_t : P \rightarrow \{0, 1\}$ are respectively the *input* and *output functions* of t , and m_0 is the *initial marking*. Given two markings m_1 and m_2 , we let $m_1 \preceq m_2$ iff $m_1(p) \leq m_2(p)$ for all $p \in P$. Given a marking m , a transition $t = (I_t, O_t)$ is *enabled* in m iff $m(p) \geq I_t(p)$ for all $p \in P$. When t is enabled in m , one can *fire* the transition t in m , which produces a new marking m' s.t. $m'(p) = m(p) - I_t(p) + O_t(p)$ for all p . This is denoted $m \xrightarrow{t} m'$, or simply $m \rightarrow m'$ when the transition identity is irrelevant. A *run* is a finite sequence $m_0 m_1 \dots m_n$ s.t. for all $1 \leq i \leq n$: $m_{i-1} \rightarrow m_i$. For a PN N , we denote by $\mathbf{Reach}(N)$ (resp. $\mathbf{Cover}(N)$) the *reachability* (*coverability*) *set* of N , i.e. the set of all markings m s.t. there exists a run $m_0 m_1 \dots m_n$ of N with $m = m_n$ ($m \preceq m_n$). The *coverability problem* asks, given a PN N and a marking m , whether $m \in \mathbf{Cover}(N)$. It is EXPSpace-complete [8]. The termination problem, i.e., whether all executions of the Petri net are finite, is decidable in EXPSpace-C [15, 16].

3 Queue-dispatch asynchronous systems

Syntax: We now define our formal model for queue-dispatch asynchronous systems. Let \mathbb{D} be a finite data domain containing an *initial value* d_0 . A *queue-*

dispatch asynchronous system (QDAS) \mathcal{A} is a tuple $\langle CQID, SQID, \Gamma, main, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$ where:

- $CQID$ and $SQID$ are respectively sets of *(c)oncurrent* and *(s)erial queues*;
- Γ is the finite set of *blocks* and $main \in \Gamma$ the *initial block*. Each block $\gamma \in \Gamma$ is a tuple $\langle S_\gamma, s_\gamma^0, f_\gamma, \Sigma, \Delta_\gamma \rangle$ where $\langle S_\gamma, s_\gamma^0, \Sigma, \Delta_\gamma \rangle$ is an LTS and $f_\gamma \in S$ a distinct final state;
- \mathcal{X} is a finite set of \mathbb{D} -valued variables;
- Σ is the set of *actions*, with $\Sigma = (\{\mathbf{dispatch}_s, \mathbf{dispatch}_a\} \times (CQID \cup SQID) \times \Gamma \setminus \{main\}) \cup \mathbf{guards}(\mathcal{X}) \cup \mathbf{assign}(\mathcal{X})$.

We assume that $SQID, CQID, \Gamma, \mathcal{X}$, and all S_γ for $\gamma \in \Gamma$ are disjoint from each other. Let $S = \bigcup_{\gamma \in \Gamma} S_\gamma$, $F = \bigcup_{\gamma \in \Gamma} \{f_\gamma\}$, $\Delta = \bigcup_{\gamma \in \Gamma} \Delta_\gamma$, and $QID = SQID \cup CQID \cup \{\iota\}$ (where $\iota \notin SQID \cup CQID$). We further assume that $\varepsilon \notin \Sigma$.

Call-task graphs: We formalize the semantics of QDAS using the notion of *call-task graph* (CTG) to describe the system's global configurations.

A configuration of a QDAS (see Fig. 2 for an example) contains a set of running tasks, represented by *task vertices* (depicted by round nodes), a set of called but unscheduled blocks, represented by *call vertices* (square nodes). Call vertices are held by queues, and the linear order of each queue is represented by *queue edges* (solid edges). Synchronous calls add an additional dependency (the caller is waiting for the termination of the callee) that is represented by a *wait edge* (dashed edges) between the caller and the callee. Wait edges are also inserted between the head of a *serial* queue and the running task that has been extracted from this queue (if it exists) to indicate that the task has to terminate before a new block can be dequeued. Note that only vertices without outgoing edges can execute a computation step, the others are currently blocked. Each node v is labeled by a block $\lambda(v)$, and by the identifier $queue(v)$ of the queue that contains it (for call vertices) or that contained it (for task vertices). Task vertices are labeled by their current state $state(v)$ (for convenience, we also label call vertices by the initial state of their respective blocks – not shown in the figure).

Example 2 The CTG in Fig. 2 depicts a configuration of a QDAS with two queues. Queue q_2 is serial (note the outgoing wait edge to the running task) and contains $\gamma_2\gamma_2\gamma_2$, and q_1 is parallel with content $\gamma_1\gamma_2$. There are 4 active tasks, two of them ($main$ and the task running γ_1) are blocked. The task running γ_3 has been dequeued from q_2 and is currently at location s . \dashv

Formally, given a QDAS $\mathcal{A} = \langle CQID, SQID, \Gamma, main, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$, a *call-task graph* over \mathcal{A} is a tuple $G_{\mathcal{A}} = \langle V, E, \lambda, queue, state \rangle$ where: $V = V_C \cup V_T$ is a finite set of *vertices*, partitioned into a set V_C of *call vertices* and a set V_T of *task vertices*; $E \subseteq V \times V$ is a set of *edges*; $\lambda : V \rightarrow \Gamma$ labels each vertex by a block; $queue : V \rightarrow QID \cup \{\iota\}$ associates each vertex to a queue identifier (or ι); and $state : V \rightarrow S$ associates each vertex to a LTS state. For




Figure 2: CTG for a QDAS with a concurrent queue q_1 and a serial queue q_2

each $q \in QID$, let $V_q = \{v \in V \mid queue(v) = q\}$. The set E is partitioned into the set E_W of *wait edges* and the set $E_Q = \bigcup_{q \in QID} E_q$ of *queue edges* where, for each $q \in QID$, $E_q = E \cap (V_q \times V_q)$.

A CTG is *empty* iff $V = \emptyset$. The *Parikh image* $\text{Parikh}(G)$ of a CTG G of \mathcal{A} is a function $f : S \rightarrow \mathbb{N}$, s.t. for all $s \in S$: $f(s) = |\{v \in V \mid state(v) = s\}|$. Given two Parikh images $\text{Parikh}(G)$ and $\text{Parikh}(G')$, we let $\text{Parikh}(G) \preceq \text{Parikh}(G')$ iff for all $s \in S$: $\text{Parikh}(G)(s) \leq \text{Parikh}(G')(s)$. A *path* (of length n) in $G_{\mathcal{A}}$ is a sequence of vertices v_0, v_1, \dots, v_n s.t. for all $1 \leq i \leq n$: $(v_{i-1}, v_i) \in E$. Such a path is *simple* iff $v_i \neq v_j$ for all $1 \leq i < j \leq n$. The *restriction* of $G_{\mathcal{A}}$ to $V' \subseteq V$ is the CTG $G'_{\mathcal{A}} = \langle V', E', \lambda', queue', state' \rangle$, where $E' = E \cap (V' \times V')$, and $\lambda', queue'$ and $state'$ are respectively the restrictions of λ , $queue$ and $state$ to V' .

In the rest of the paper, we assume that all the CTG we consider are *well-formed*, i.e., they fulfill the following requirements:

1. For each $v \in V_T$: $state(v) \in S_{\lambda(v)}$ where $S_{\lambda(v)}$ are the states of $\mathcal{TS}_{\lambda(v)}$.
2. Each *call* vertex has at most one outgoing (queue or wait) edge, at most one incoming *wait* edge, and at most one incoming *queue* edge. Each *task* vertex has at most one outgoing, and at most one incoming *wait* edge.
3. For each $q \in QID$, the restriction of $G_{\mathcal{A}}$ to V_q is either empty or contains one and only one simple path of length $|V_q| - 1$. Intuitively, this ensures the well-formedness of the queues.
4. For each $q \in SQID$, there is at most one task vertex v s.t. $queue(v) = q$. This ensures that queues in $SQID$ indeed force the serial execution of its members.

For convenience, we also introduce the following notations. Let $G_{\mathcal{A}}$ be a CTG, and let q be a queue identifier of \mathcal{A} . Then, $head(q, G_{\mathcal{A}})$ and $tail(q, G_{\mathcal{A}})$ denote respectively the head and the tail of q in the configuration described by $G_{\mathcal{A}}$, that is, $head(q, G_{\mathcal{A}})$ is the call vertex $v \in V_q$ that has no incoming queue edge, or \perp , if such a vertex does not exist; and $tail(q, G_{\mathcal{A}})$ is the call vertex $v \in V_q$ that has no outgoing *queue* edge (but possibly an outgoing *wait* edge), or \perp , if such a vertex does not exist. Remark that, when they exist, these vertices are necessarily unique because of the well-formedness assumptions. Finally, we say that a vertex v is *unblocked* iff it has no outgoing edge, and that it is *final* iff (i) v is an *unblocked task* vertex and (ii) $state(v) = f_{\lambda}(v)$ (that is, v represents a task that has reached the final state of its transition system and is not waiting on another task).

Let us now define several operations on CTG. We will rely on these operations when defining the formal semantics of QDAS. Let \mathcal{A} be a QDAS and $G_{\mathcal{A}} = \langle V, E, \lambda, queue, state \rangle$ be a CTG for \mathcal{A} . Then:

- for all $v \in V$: $G \setminus v$ is the restriction of G to $V \setminus \{v\}$.
- for all $\gamma \in \Gamma$ and $q \in QID$, $enqueue(q, \gamma)(G_{\mathcal{A}})$ is the CTG $\langle V', E', \lambda', queue', state \rangle$ where: $V' = V \cup \{v'\}$, v' is a fresh queue vertex, $\lambda(v') = \gamma$, $queue(v') = q$, $state(v') = s_{\gamma}^0$, and for all $v \in V$: $\lambda'(v) = \lambda(v)$ and $queue'(v) = queue(v)$. Finally, $E' = E \cup E_1 \cup E_2$, where: (i) $E_1 = \{(v', tail(G_{\mathcal{A}}, q))\}$ if $tail(G_{\mathcal{A}}, q) \neq \perp$, and $E_1 = \emptyset$ otherwise, and (ii) if $v \in V$ is a *task* node s.t. $queue(v) = q \in SQID$, then $E_2 = \{(v', v)\}$, otherwise $E_2 = \emptyset$. Intuitively, this operation inserts a call to γ in the queue q , by creating a new vertex v' and adding an edge to maintain the FIFO ordering, if necessary (set E_1). In the case of a *serial* queue that was empty before the enqueue, a supplementary edge (in set E_2) might be necessary to ensure that v' is blocked by a currently running v which has been extracted from q .
- for all $q \in QID$, if $head(q)$ is different from \perp and *unblocked*, then $dequeue(q)(G_{\mathcal{A}})$ is the CTG $\langle V'_C \cup V'_T, E', \lambda, queue, state \rangle$ where $V'_C = V_C \setminus \{head(q)\}$ and $V'_T = V_T \cup \{head(q)\}$. Otherwise, $head(q) = \perp$ and $dequeue(q)(G_{\mathcal{A}})$ is undefined. Intuitively, this operation removes the first (with respect to the FIFO ordering) block from q and turns the corresponding *call* vertex $head(q)$ into a *task* vertex, meaning that the block is now running as a task.
- for all $\delta = (s, a, s') \in \Delta$, $step(\delta)(G_{\mathcal{A}})$ is a *set* of CTG defined as follows. $\langle V, E, \lambda, queue, state' \rangle \in step(\delta)(G_{\mathcal{A}})$ iff there exists an *unblocked* $v \in V_T$ s.t. $state(v) = s$, $state'(v) = s'$ and for all $v' \neq v$: $state'(v') = state(v')$. Remark that $step(\delta)(G_{\mathcal{A}})$ can be empty. Intuitively, each graph in $step(\delta)(G_{\mathcal{A}})$ corresponds to the firing of an a -labeled transition by a task that is not blocked.
- for all unblocked $v \in V \cup \{\perp\}$, all $v' \in V$: $letwait(v, v')(G_{\mathcal{A}})$ is either the CTG $G_{\mathcal{A}}$ if $v = \perp$, or the CTG $\langle V, E \cup (v, v'), \lambda, queue, state \rangle$ if $v \neq \perp$. Intuitively, this operation adds a wait edge between nodes v and v' when $v \neq \perp$, and does not modify the CTG otherwise.

Semantics of Qdas: For a QDAS \mathcal{A} with set of variables \mathcal{X} , a *configuration* is a pair (G, \vec{d}) , where G is a CTG of \mathcal{A} and $\vec{d} \in \text{vals}(\mathcal{X})$. The operational semantics of \mathcal{A} is given as a transition system $\llbracket \mathcal{A} \rrbracket$ whose states are configurations of \mathcal{A} ; and whose transitions reflect the semantics of the actions labeling the transitions of the QDAS. Formally, given a QDAS $\mathcal{A} = \langle CQID, SQID, \Gamma, main, \mathcal{X}, \Sigma, (\mathcal{TS}_{\gamma})_{\gamma \in \Gamma} \rangle$, $\llbracket \mathcal{A} \rrbracket$ is the labeled transition system $\langle C, c^0, \tilde{\Sigma}, \Longrightarrow \rangle$ where: (i) C contains all the pairs (G, \vec{d}) where $\vec{d} \in \text{vals}(\mathcal{X})$, and G is a CTG of \mathcal{A} , (ii) $c^0 = (G^0, \vec{d}^0)$ with $\vec{d}^0(x) = d_0$ for all $x \in \mathcal{X}$, and $G^0 = \langle \{v^0\}, \emptyset, \lambda, queue, state \rangle$, where v^0 is a *task node*, $\lambda(v^0) = main$, $state(v^0) = s_{main}^0$ and $queue(v^0) = \perp$, (iii) $\tilde{\Sigma} = \Sigma \cup \{\varepsilon\}$ and (iv) $((G, \vec{d}), a, (G', \vec{d}')) \in \Longrightarrow$ iff one of the following holds:

Async. dispatch: $a = \text{dispatch}_a(q, \gamma)$, $\vec{d}' = \vec{d}$, and there are $\delta = (s, a, s') \in \Delta$ and $G'' \in \text{step}(\delta)(G)$ s.t.: $G' = \text{enqueue}(q, \gamma)(G'')$.

Sync. dispatch: $a = \text{dispatch}_s(q, \gamma)$, $\vec{d}' = \vec{d}$ and there are $\delta = (s, a, s') \in \Delta$ and $G'' \in \text{step}(\delta)(G)$ s.t.: $G' = \text{letwait}(v, v')(\text{enqueue}(q, \gamma)(G''))$ where v is the node whose *state* has changed during the *step* operation, and v' is the fresh node that has been created by the *enqueue* operation. That is, a queue vertex v' labeled by γ is added to q and a *wait* edge is added between the node v representing the task that performs the *synchronous* dispatch, and v' , as the dispatch is *synchronous*.

Test: $a = g \in \text{guards}(\mathcal{X})$, $\vec{d}' = \vec{d}$, $\vec{d} \models g$, and there is $\delta = (s, a, s') \in \Delta$ s.t. $G' \in \text{step}(\delta)(G)$.

Assignment: $a = x \leftarrow v \in \text{assign}(\mathcal{X})$, $\vec{d}'(x) = v$, for all $x' \neq x$: $\vec{d}'(x) = \vec{d}(x)$ and there is $\delta = (s, a, s') \in \Delta$ s.t. $G' \in \text{step}(\delta)(G)$.

Scheduler action: $a = \varepsilon$, $\vec{d}' = \vec{d}$ and:

- either there is a final vertex v s.t. $G' = G \setminus v$;
- or there is $q \in CQID$ s.t. $\text{head}(q, G) \neq \perp$ and $G' = \text{dequeue}(q)(G)$. That is, the scheduler schedules a block (represented by v) from a concurrent queue.
- or there is $q \in SQID$ s.t. $\text{head}(q, G) = v$, v is *unblocked*, as well as $G' = \text{letwait}(\text{head}(q, G''), v)(G'')$ and $G'' = \text{dequeue}(q)(G)$. That is, the scheduler schedules a block (represented by v) from the serial queue q . As the queue is serial, a *wait* edge is inserted between the next waiting block in q (now represented by $\text{head}(q, G'')$) and v .

A run ρ of a QDAS is an alternating sequence $c_0 a_1 c_1 a_2 \dots a_n c_n$ of configurations and actions where $(c_i, a_{i+1}, c_{i+1}) \in \text{transitions}$ for all $0 \leq i < n$ and $c_0 = c^0$. A run is *finite* if this sequence is finite. A configuration c is *reachable* in \mathcal{A} iff there exists a finite run $c_0 a_1 c_1 a_2 \dots a_n c_n$ of \mathcal{A} s.t. $c_n = c$. We denote by $\text{Reach}(\mathcal{A})$ the set of all reachable configurations of \mathcal{A} .

The decision problem on QDAS we mainly consider in this work is the *Parikh coverability problem*: given a QDAS \mathcal{A} with set of locations S and a function $f : S \mapsto \mathbb{N}$, it asks whether there is $c = (G, \vec{d}) \in \text{Reach}(\mathcal{A})$ s.t. $f \preceq \text{Parikh}(G)$. When the answer to this question is ‘yes’, we say that f is *Parikh-coverable* in \mathcal{A} . It is well-known that meaningful verification questions can be reduced to this problem. For instance, consider a *mutual exclusion* question, asking whether it is possible to reach, in a QDAS \mathcal{A} , a configuration in which at least two tasks are executing the same block γ and are in the same control state s . If yes, the mutual exclusion (of control state s) is violated. This can be encoded into an instance of the Parikh coverability problem, where $f(s) = 2$ and $f(s') = 0$ for all $s' \neq s$, and would allow, for example, to verify if there are more than one block of type **increase** running in Example 1.




Figure 3: The two possible forms of reachable CTGs in a synchronous QDAS

In addition, we look at the *(universal) termination problem*: given a QDAS \mathcal{A} , it asks whether all executions of \mathcal{A} are finite, i.e., there is no infinite run of \mathcal{A} . Regarding Example 1, this permits to test whether the `main` task terminates, i.e., all dispatched blocks terminate.

4 From the Parikh coverability problem to Termination

Before regarding the termination problem, we first study in this section the Parikh coverability problem from a computational point of view. As expected, this problem is undecidable in general. However, when restricting the types of queues and dispatches that are allowed, it is possible to retain decidability. In these cases, we characterize the complexity of the problem. Formally, we consider the following subclasses of QDAS. A QDAS \mathcal{A} with set of transitions Δ , set of serial queues $SQID$ and set of concurrent queues $CQID$, is *synchronous* iff there exists no $(s, a, s') \in \Delta$ with $a \in \{\text{dispatch_a}\} \times QID \times \Gamma$; it is *asynchronous* iff there exists no $(s, a, s') \in \Delta$ with $a \in \{\text{dispatch_s}\} \times QID \times \Gamma$; it is *concurrent* iff $SQID = \emptyset$ and $CQID \neq \emptyset$; it is *serial* iff $CQID = \emptyset$ and $SQID \neq \emptyset$; it is *queueless* iff $CQID = SQID = \emptyset$.

Queueless Qdas: In a queueless QDAS, there is no dispatch possible, so the only task that can execute at all time is the `main` one. Thus, configurations of queueless QDAS can be encoded as tuples (s, \vec{d}) , where s is a state of `main`, and \vec{d} is a valuation of the variables. Hence queueless QDAS are essentially LTS with variables over a finite data domain, thus:

Proposition 2 *The Parikh coverability is PSPACE-C for queueless QDAS.*

Synchronous Qdas: In synchronous QDAS, there is no concurrency in the sense there is at most one running task that can fire an action at all times. All the other tasks have necessarily performed a *synchronous* dispatch and are thus blocked. More precisely, in every reachable configuration (G, \vec{d}) of a synchronous QDAS, G is of one of the forms depicted in Fig. 3 (i.e. $v_0, \dots, v_{n-1} \in V_T$ and either $v_n \in V_T$ or $v_n \in V_C$). When the current CTG is of the form Fig. 3(a), the only possible action is that the scheduler starts running v_n 's block and we obtain a graph of the form Fig. 3(b). In the case where the CTG is of the form (a), either v_n terminates, which removes v_n from the CTG, or v_n executes an internal action, which does not change the shape of the CTG, or v_n does a synchronous call, which adds a call vertex as successor of v_n which will be

directly scheduled. W.l.o.g., we assume in the following that for synchronous QDAS the combined action of dispatch_s and scheduling the dispatched block is atomic.

For a CTG G and $w \in S^*$, we write $G \triangleright w$ iff for all $0 \leq i \leq n$: $w_i = \text{state}(v_i)$ and the empty CTG is mapped to the empty word ε . Given a synchronous QDAS \mathcal{A} with set of local states S as before, we can build a pushdown system with data $\mathcal{P}_{\mathcal{A}}$ such that, at all times, the current location of $\mathcal{P}_{\mathcal{A}}$ encodes the current location of the (single) running block in \mathcal{A} , and the stack content records the sequence of synchronous dispatches, as described above. A guard or assignment in \mathcal{A} is kept as is in $\mathcal{P}_{\mathcal{A}}$. A synchronous dispatch $(s, \text{dispatch}_s(q, \gamma), s')$ in \mathcal{A} is simulated by a push of s' (to record the local state that has to be reached when the callee terminates) and moves the current state of $\mathcal{P}_{\mathcal{A}}$ to the initial state of γ . The termination of a block is simulated by a pop (and we encode the termination of *main* in testing the stack's emptiness).

Proposition 3 *Given a synchronous QDAS \mathcal{A} , then we can construct a pushdown system with data $\mathcal{P}_{\mathcal{A}}$ such that the following holds: for any run $\rho = c_0 a_1 c_1 \dots a_n c_n$ of \mathcal{A} , there exists a run $\pi = x_0 a_1 x_1 \dots a_n x_n$ in $\mathcal{P}_{\mathcal{A}}$ such that for all $c_i = (G_i, \vec{d}_i)$ and $x_i = (s_i, w_i, \vec{d}_i)$ we have $\vec{d}_i = \vec{d}_i$ and $G_i \triangleright w_i$ ($0 \leq i \leq n$), and vice versa.*

The previous proposition allows to derive results on the reachability problem. However, we are interested in the Parikh coverability problem. Let f be a Parikh image of \mathcal{A} . Then, by Proposition 4, looking for a reachable configuration of \mathcal{A} that covers f amounts to finding a reachable configuration (s_i, w_i, \vec{d}_i) of $\mathcal{P}_{\mathcal{A}}$ s.t. the Parikh image P of w_i is s.t. $f \preceq P$ (as the CTG is encoded by the stack content w_i). To achieve this, we augment $\mathcal{P}_{\mathcal{A}}$ with a *widget* that works as follows. In any location of $\mathcal{P}_{\mathcal{A}}$, we can jump non-deterministically to the widget. Then, the widget pops all the values from the stack, and checks that at least $f(s)$ symbols s are present on the stack. The widget jumps to an accepting state iff it is the case. We call $\mathcal{P}_{\mathcal{A},f}$ the resulting PDS. Clearly, one can build such a widget for all f , and this effectively reduces the Parikh coverability problem of QDAS to the location reachability problem of PDS. Moreover, for all f , the widget is of size exponential in $|S|$ and exponential in the binary encoding of $\max_{s \in S} f(s)$. Hence, building $\mathcal{P}_{\mathcal{A},f}$ requires exponential time:

Proposition 4 *Given a synchronous QDAS \mathcal{A} with states S and a function $f : S \rightarrow \mathbb{N}$, then one can generate a PDS $\mathcal{P}_{\mathcal{A},f}$ of size exponential in \mathcal{A} and a state s of $\mathcal{P}_{\mathcal{A},f}$, s.t. $\mathcal{P}_{\mathcal{A},f}$ reaches s iff f is Parikh coverable in \mathcal{A} .*

As testing emptiness of a pushdown system without data is PTIME-C [4], the Parikh coverability problem is in EXPTIME for *synchronous* QDAS (with both types of queues). A matching lower bound is obtained by reducing the reachability question of PDS with data (see Proposition 1). This reduction requires only one *concurrent* queue, so the Parikh reachability problem is EXPTIME-hard for *synchronous concurrent* QDAS. Hence we derive the following:

Theorem 1 *The Parikh coverability problem is EXPTIME-C for synchronous and for synchronous concurrent QDAS.*

Let us take a closer look on the dispatches that happen in runs of synchronous QDAS that have only *serial* queues. Here, each task except the **main** task blocks the queue it is started from. Hence, any other block dispatched to these already blocked queues deadlocks. Thus, all reachable CTG have at most $|SQID| + 2$ vertices. Hence, the pushdown systems used in all previous constructions have bounded stack height, and we can apply test on a finite transition system. The lower bound can be derived from Proposition 2. by testing the emptiness of the intersection of n finite processes, that is PSPACE-complete [14].

Theorem 2 *The Parikh coverability problem is PSPACE-C for serial synchronous QDAS.*

Concurrent asynchronous Qdas: Let us now establish a relationship between *concurrent asynchronous* QDAS and Petri nets that proves that the Parikh coverability problem is EXPSPACE-complete. We first show how to reduce the QDAS Parikh coverability problem to the Petri net coverability problem. Given a concurrent asynchronous QDAS \mathcal{A} , we construct a Petri net $N_{\mathcal{A}}$ as follows: The places of $N_{\mathcal{A}}$ are $(\mathcal{X} \times \mathbb{D}) \cup S$. Each place $s \in S$ counts how many blocks are currently running and are in state s . Each place (x, d) encodes the fact that variable x contains value d in the current valuation. Remark that we have no place to encode the contents of the queue, as the dispatch of block γ directly creates a new token in s_{γ}^0 . This encoding is, however, correct with respect to the *Parikh coverability problem*, as $\text{Parikh}(G)$ does not distinguish between a block γ that is waiting in a queue, and a task executing γ in its initial state. Thus:

Proposition 5 *For all concurrent asynchronous QDAS \mathcal{A} with set of location S , we can build, in polynomial time, a Petri net $N_{\mathcal{A}}$ s.t. f is Parikh-coverable in \mathcal{A} iff $m \in \text{Cover}(N_{\mathcal{A}})$, where m is the marking s.t. for all $s \in S$: $m(s) = f(s)$ and for all $p \in P \setminus S$: $m(p) = 0$.*

Let us now reduce the Petri net coverability problem to the QDAS Parikh coverability problem. Let $N = \langle P, T, m_0 \rangle$ be a Petri net. We associate to N the concurrent asynchronous QDAS $\mathcal{A}_N = \langle CQID, \emptyset, \Gamma, \mathbf{main}, \mathcal{X}, \Sigma, (\mathcal{TS}_{\gamma})_{\gamma \in \Gamma} \rangle$, on the finite domain $\mathbb{D} = \{0, 1\}$, where $CQID = \{C\}$, $\Gamma = \{\mathbf{main}, \mathbf{trans}\} \cup P$, $\mathcal{X} = \{v_p \mid p \in P\}$ and $(\mathcal{TS}_{\gamma})_{\gamma \in \Gamma}$ is given by the pseudo-code in Fig. 4 (this construction is an extension of a construction found in [10]). We assume that, for $\gamma \in \{\mathbf{trans}, \mathbf{main}\}$ s_{γ}^{ℓ} is the location of γ 's LTS that is reached when the control reaches line ℓ . Let $G = \langle V, E, \lambda, \text{queue}, \text{state} \rangle$ be a CTG for \mathcal{A}_N , and let m be a marking of N . Then, we say that G *encodes* m , written $G \triangleright m$ iff (i) $\text{Parikh}(G)(s_{\mathbf{trans}}^{14}) = \text{Parikh}(G)(s_{\mathbf{main}}^8) = 1$, (ii) for all $p \in P$: $\text{Parikh}(G)(s_p^0) = m(p)$ and (iii) for all $p \in P$, for all $s \in S_p \setminus \{s_p^0\}$: $\text{Parikh}(G)(s) = 0$. Thus, intuitively, a CTG G encodes a marking m iff **main** is at line 8, **trans** is at line

```

1 def main():
2   for each  $p \in P$ :
3      $v_p := 0$ 
4     select  $k_p \in \{0, \dots, m_0(p)\}$ 
5     for  $i = 0 \dots k_p$ :
6       dispatch_a( $C, p()$ )
7     dispatch_a( $C, trans()$ )
8     while(true): do nothing
9
10  block p(): // For all  $p \in P$ 
11    while( $v_p = 0$ ): do nothing
12     $v_p := 0$ 
13
12 block trans():
13   while(true):
14     select  $t = (I_t, O_t) \in T$ 
15     for each  $p \in P$  s.t.  $I_t(p) = 1$ :
16        $v_p := true$ 
17     while( $\exists p \in P: v_p = 1$ ): do nothing
18     for each  $p \in P$  s.t.  $O_t(p) = 1$ :
19       dispatch_a( $C, p()$ )

```

Figure 4: Encoding of Petri net coverability $\langle P, T, m_0 \rangle$ by a QDAS

14, $m(p)$ counts the number of p blocks that are either in C or executing but at their initial state, and there are no p blocks that are in state s_p^{mid} or s_p^{fin} .

The intuition behind the construction is as follows. Each run of the QDAS \mathcal{A}_N starts with an initialization phase, where **main** initializes all the v_p variables to 0 and dispatches, for all $p \in P$, k_p blocks p with $k_p \leq m_0(p)$, then dispatches a call to **trans**. At that point, the only possible action is that the scheduler dequeues all the blocks. All the p tasks are then blocked, as they need that $v_p = 1$ to proceed and terminate. Then, **trans** cyclically picks a transition t , sets to 1 all the variables v_p s.t. t consumes a token in p , and waits that all the v_p variables return to 0. This can only happen because *at least* $I_t(p)$ p tasks have terminated, for all $p \in P$. So, when **trans** reaches line 19, the encoded marking has been decreased by *at least* I_t . Remark that more than $I_t(p)$ p tasks could terminate, as they run concurrently, and the lines 11 and 12 do not execute atomically. Then, **trans** dispatches one new p block iff t produces a token in p . This increases the encoded marking by O_t , so the effect of one iteration of the main **while** loop of **trans** is to simulate the effect of t , plus a possible token loss. Hence, the resulting marking is guaranteed to be in $Cover(N)$ (but maybe not in $Reach(N)$). This is formalized by the following proposition:

Proposition 6 *For all Petri nets N , we can build, in polynomial time, a concurrent asynchronous QDAS \mathcal{A}_N s.t. $m \in Cover(N)$ iff there exists $(G, \vec{d}) \in Reach(\mathcal{A}_N)$ with $G \triangleright m$.*

Theorem 3 *The Parikh coverability problem is EXPSpace-complete for concurrent asynchronous QDAS.*

Asynchronous Serial Qdas: Let us show that for the class of QDAS with one serial queue, and where asynchronous dispatches are allowed, the Parikh coverability problem is *undecidable*. We establish this by a reduction from the control-state reachability problem in a fifo system which is known to be undecidable [5].

Intuitively, we use the serial queue to model the unbounded, reliable fifo queue where sending a message m is encoded as asynchronously dispatching a block γ_m . This block γ_m contains the control-flow of receiving m , i.e., that will

resume the fifo system's execution directly after receiving m . The fifo system's global state is guarded in a global variable. Receiving a certain message m is encoded as terminating the currently running task and assuring (via a global variable) that the succeeding task's type is the one of the expected message.

Theorem 4 *The Parikh coverability problem is undecidable for asynchronous QDAS with at least one serial queue.*

Concurrent Qdas: Let us show that, once we allow both synchronous and asynchronous dispatches in a *concurrent* QDAS, the Parikh coverability problem becomes undecidable. For that purpose, we reduce the reachability problem of two counter systems.

The crux of the construction is the use of variables, i.e., global memory, to implement a rendez-vous synchronization. Given two distinct tasks, one can use their nested access to two lock variables to guard a shared data variable by assuring that a value written to the variable must be read before it is overwritten.

Let us give the construction's intuition: Each counter is encoded similarly to the construction for synchronous QDAS as pushdown stack over a singleton alphabet, i.e., a sequence of nested synchronous dispatched blocks, these are controlled via rendez-vous from the *main* task that in the beginning asynchronously dispatched the two counters.

Theorem 5 *The Parikh coverability problem is undecidable for concurrent QDAS that use both synchronous and asynchronous dispatches.*

Termination Problem: We use the previous constructions to directly lift the undecidability results from the Parikh coverability problem to the termination problem. The close connection of synchronous QDAS with PDS (with data) allows to directly derive an EXPTIME algorithm for the termination problem from the emptiness testing of Büchi PDS [9]. Up to our knowledge, no completeness result is known for the latter problem, thus leaving a gap to the directly derivable PSPACE-hardness via finite systems. The result for asynchronous concurrent QDAS directly follows from Petri nets [15, 16].

Theorem 6 *The termination problem is PSPACE-C for synchronous serial QDAS, it is in EXPTIME and PSPACE-hard for synchronous QDAS, and it is EXPSpace-C for asynchronous concurrent QDAS. It is undecidable for asynchronous serial QDAS, and QDAS that use both synchronous and asynchronous dispatches.*

5 Extending QDAS with Fork/Join

We return to the introductory matrix multiplication example. The crux of the algorithm is the parallel for-loop that *forks* a finite number of subtasks and waits for their termination (*join*). The latter had to be implemented via a global semaphore which (*i*) restricts the number of forkable tasks by the underlying

finite value domain, and (ii) needs to be properly guarded by the programmer for access outside fork and join. In the following we thus want to extend QDAS by an explicit fork/join construct (which also exists in GCD). Further, the given matrix multiplication algorithm depended on an a priori fix size for the factor matrices, however, in practice, one wants to verify the algorithm for any possible (correct) input of any size. Thus, we need to consider the verification of extended QDAS where the number of forked tasks is parametrized by the input.

As fork/join behaviour relies on asynchronously dispatching tasks on a concurrent queue, we ignore in the following synchronous dispatches and serial queues, thus also partially avoiding the previous basic undecidability results. Note that asynchronous concurrent QDAS can be regarded as over-approximations of all other classes of QDAS.

QDAS extended by fork/join An QDAS *extended by fork/join* (EQDAS) is a tuple $\langle CQID, \emptyset, \Gamma, main, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$ that is equivalent to a QDAS except that we replace in Σ the synchronous dispatch by the following action: $\{\text{forkjoin}\} \times CQID \times \Gamma \times (\mathbb{N} \cup \{*\})$. The *parameter* of a **forkjoin** action is the last value of the tuple. An EQDAS is **-free* if in all \mathcal{TS}_γ for $\gamma \in \Gamma$ the parameter of the **forkjoin** action is not $*$.

The semantics of an EQDAS is given analogous to standard QDAS as transition system $\langle C, c^0, \Sigma, \Rightarrow \rangle$ where we additionally extend the transition relation \Rightarrow given by tuples $((G, \vec{d}), a, (G', \vec{d}'))$ by the following case:

Fork/join: $a = \text{forkjoin}(q, \gamma, p)$ with $p \in (\mathbb{N} \cup \{*\})$, $\vec{d}' = \vec{d}$ and there are $\delta = (s, a, s') \in \Delta$, and $G'' \in \text{step}(\delta(G))$ such that: if $p = *$ then we choose non-deterministically an $n \in \mathbb{N}$, else $n = p$, so that $G' = G''_n$ where $G''_0 = G''$ and for $0 < i \leq n$ we define $G''_{i+1} = \text{letwait}(v, v'_{i+1})(\text{enqueue}(q, \gamma_{i+1})(G''_i))$ where v is the node whose *state* has changed during the **step** operation, and v'_{i+1} is the fresh node that has been created by the **enqueue** operation.

Intuitively, a **forkjoin** action appends a sequence of blocks to a queue by additionally adding a wait edge to each newly create node. Hence, the join is modeled by a separate action that is taken by the scheduler after deleting the wait edges.

The *extended Parikh coverability problem* asks, given an EQDAS \mathcal{A} with locations \mathcal{S} and a mapping $f : \mathcal{S} \rightarrow \mathbb{N}$, whether there exists $c = (G, \vec{d}) \in \text{Reach}(\mathcal{A})$ with $f \preceq \text{Parikh}(G)$. The *extended termination problem* asks, given an EQDAS \mathcal{A} whether there is no infinite run possible in \mathcal{A} .

As **forkjoin** actions with parameter 1 are semantically equivalent to a synchronous dispatch action, we can directly reduce the two counter machine simulation from the proof of Theorem 5 to EQDAS.

Theorem 7 *Both the extended Parikh coverability and extended termination problem are undecidable.*

Consequently, we focus on two distinct over-approximations for EQDAS in the following that allow us to give approximative answers to our verification problems.

***-free eQdas:** Given an EQDAS \mathcal{A} that is *-free. We construct a Petri net $N_{\mathcal{A}}^{\times}$ by extending the previous construction from asynchronous concurrent QDAS to Petri nets as follows: As in the EQDAS semantics we split a single **forkjoin** action of a block γ on a queue q with parameter $n \in \mathbb{N}$ into (i) a fork transition that creates n new tokens in s_{γ}^0 , and (ii) a subsequent join transition that depends on taking n tokens from the place representing f_{γ} . Analogous to the proof of Proposition 5 we can show the following:

Proposition 7 *For all *-free EQDAS with set of location \mathcal{S} , we can build in polynomial time a Petri net $N_{\mathcal{A}}^{\times}$ st. f is Parikh-coverable in \mathcal{A} if $m \in \text{Cover}(N_{\mathcal{A}}^{\times})$, where m is the marking s.t. for all $s \in S$: $m(s) = f(s)$ and for all $p \in P \setminus S$: $m(p) = 0$. Further, if $N_{\mathcal{A}}^{\times}$ terminates, then \mathcal{A} is guaranteed to terminate.*

As coverability and termination are decidable for Petri nets, we can decide extended Parikh coverability and extended termination on this over-abstraction.

eQdas with * parametrized fork/join: Given an EQDAS \mathcal{A} that is not *-free, we construct a Petri net $N_{\mathcal{A}}^*$ as follows starting from the construction for asynchronous concurrent QDAS: For **forkjoin** actions whose parameter is not *, we proceed as in the above construction for *-free EQDAS. However, we need to model the forking of an arbitrary number of blocks when the parameter of the **forkjoin** action equals *. For this, we use Petri nets extended with ω -arcs. An outgoing arc of a transition labeled with ω adds an arbitrary number of tokens to the corresponding place, thus, we translate the fork of block γ into an ω -transition leading to place s_{γ}^0 . The join is approximated by a transition that non-deterministically chose to advance the original workflow, ignoring not already terminated forked tasks. Thus by extending the proof of Proposition 5:

Proposition 8 *For all EQDAS with set of location \mathcal{S} , we can build in polynomial time a Petri net $N_{\mathcal{A}}^*$ st. f is Parikh-coverable in \mathcal{A} if $m \in \text{Cover}(N_{\mathcal{A}}^*)$, where m is the marking s.t. for all $s \in S$: $m(s) = f(s)$ and for all $p \in P \setminus S$: $m(p) = 0$. Further, if $N_{\mathcal{A}}^*$ terminates, then \mathcal{A} is guaranteed to terminate.*

We have recently shown that the termination problem is decidable for Petri nets with ω -arcs [12]. Hence, also extended termination is decidable on the previous abstraction.

With respect to coverability, we can replace the ω -arcs of $N_{\mathcal{A}}^*$ by a non-deterministic loop that adds an arbitrary number of tokens to the original arc's target place. Note that this simple trick does not work for verifying termination. Consequently, we can use the known algorithms for coverability on this polynomially larger standard Petri net, and hence the extended Parikh coverability problem is decidable on this abstraction.

6 Conclusion & Outlook

We introduce the, up to our knowledge, first formal model that grasps the core of GCD, and that allows to derive basic results on the decidability of verification

question thereupon. Due to the obvious undecidability issues of the model, we currently focus on several under- and over-approximative approaches (e.g., language bounded verification, graph minor based abstractions, novel Petri net extensions [12]) as well as enhancements for additional GCD features like task groups, priorities, and timer events.

References

- [1] Grand Central Dispatch (GCD) Reference. Technical report, Apple Inc., 2010.
- [2] Concurrency Programming Guide. Technical report, Apple Inc., 2011.
- [3] A. Bouajjani and M. Emmi. Analysis of recursively parallel programs. In *Proc. of POPL’12*, p.203–214, 2012.
- [4] A. Bouajjani et al. Reachability analysis of pushdown automata: Application to model-checking. In *Proc. of CONCUR’97, LNCS 1243*, p.135–150. Springer, 1997.
- [5] D. Brand and P. Zafiropulo. On Communicating Finite-State Machines. *Journal of the ACM*, 30(2):323–342, 1983.
- [6] R. Chadha and M. Viswanathan. Decidability results for well-structured transition systems with auxiliary storage. In *Proc. of CONCUR’07, LNCS 4703*, pages 136–150, 2007.
- [7] R. Chadha and M. Viswanathan. Deciding branching time properties for asynchronous programs. *Theoretical Computer Science*, 410(42):4169–4179, 2009.
- [8] J. Esparza. Decidability and complexity of Petri net problems — an introduction. In *Lectures on Petri nets I, LNCS 1491*. Springer, 1998.
- [9] J. Esparza et al. Efficient algorithms for model checking pushdown systems. In *Proc. of CAV’00, LNCS 1855*, pages 232–247. Springer, 2000.
- [10] P. Ganty and R. Majumdar. Algorithmic verification of asynchronous programs. *TOPLAS*, 34(1), 2012.
- [11] P. Ganty, R. Majumdar, and A. Rybalchenko. Verifying liveness for asynchronous programs. In *Proc. of POPL’09*, p.102–113. ACM Press, 2009.
- [12] G. Geeraerts et al. ω -petri nets. ULB Research Report. <http://www.ulb.ac.be/di/verif/ggeeraer/papers/wPetri.pdf>.
- [13] R. Jhala and R. Majumdar. Interprocedural analysis of asynchronous programs. In *Proc. of POPL’07*, p.339–350. ACM Press, 2007.

- [14] D. Kozen. Lower bounds for natural proof systems. In *Proc. of FOCS'77*, p.254–266. IEEE Comp. Soc. Press, 1977.
- [15] R. Lipton. The Reachability Problem Requires Exponential Space. Techreport 62, Yale University, 1976
- [16] C. Rackoff. The Covering and Boundedness Problem for Vector Addition Systems. *TCS*, 6:223–231, 1978.
- [17] K. Sen and M. Viswanathan. Model checking multithreaded programs with asynchronous atomic methods. In *Proc. of CAV'06, LNCS 4144*, pages 300–314, 2006.

A Proof for Section 2

Proposition 1 *The reachability problem is EXPTIME-C for PDS with data.*

Proof. For the upper bound, we generate a reachability-equivalent PDS (without data) by encoding all possible data valuations into the pushdown system's states. This leads to an exponential blowup of the state space. The lower bound can be derived from the reduction of the emptiness test of the intersection of a context-free language with n regular languages that is known to be EXPTIME-hard (hardness follows easily by a reduction from linearly bounded alternating Turing machines; a closely related problem, the reachability of pushdown systems with checkpoints, is shown to be EXPTIME-hard in (*)).

(*) Javier Esparza, Antonín Kučera, and Stefan Schwoon: *Model checking LTL with regular valuations for pushdown systems*, in Information and Computation, 186(2):355–376, 2003.

B Proofs of Section 4

Synchronous Qdas: Let \mathcal{A} be a synchronous QDAS with a set of locations S , a set of rules Δ , a set of final states F , and set of queues $SQID$. Let G be a CTG of one of the forms given in Fig. 3, and let $w = w_0w_1 \cdots w_n$ be a word in S^* . Then, G is encoded by w , written $G \triangleright w$, iff for all $0 \leq i \leq n$: $w_i = state(v_i)$ and the empty CTG is mapped to the empty word ε .

Given a synchronous QDAS $\mathcal{A} = \langle CQID, SQID, \Gamma, main, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$ with set of local states S as before, we build a pushdown system with data $\mathcal{P}_\mathcal{A} = \langle Y, \mathcal{X}, y^0, S, \Sigma_\mathcal{P}, \Delta_\mathcal{P} \rangle$ where:

- the set of states is $Y = S \cup \{\varepsilon\}$ and the initial state is $y^0 = s_{main}^0$
- $\Sigma_\mathcal{P} = (\{\text{push}, \text{pop}\} \times S) \cup \{\text{empty?}\} \cup \text{guards}(\mathcal{X}) \cup \text{assign}(\mathcal{X})$
- a tuple (y, a, y') is a transition rule in $\Delta_\mathcal{P} \subseteq Y \times \Sigma_\mathcal{P} \times Y$ iff
 - $\sim a \in \text{guards}(\mathcal{X}) \cup \text{assign}(\mathcal{X})$ and $(y, a, y') \in \Delta$
 - $\sim a = \text{push}(s'), (s, \text{dispatch}_s(q, \gamma), s') \in \Delta$ and $y' = s_\gamma^0$
 - $\sim a = \text{pop}(s), y \in F$ and $y' = s$
 - $\sim a = \text{empty?}, y = f_{main}$, and $y' = \varepsilon$.

Thus, at all times, the current location of $P_\mathcal{A}$ encodes the current location of the (single) running block in \mathcal{A} , and the stack content records the sequence of synchronous dispatches, as described above. A guard or assignment in \mathcal{A} is kept as is in $P_\mathcal{A}$. A synchronous dispatch $(s, \text{dispatch}_s(q, \gamma), s')$ in \mathcal{A} is simulated by a push of s' (to record the local state that has to be reached when the callee terminates) and moves the current state of $P_\mathcal{A}$ to the initial state of γ . The termination of a block is simulated by a pop (and we use the **empty?** action for the termination of *main*).

Proposition 3 *Given a synchronous QDAS \mathcal{A} , then we can construct a pushdown system with data $\mathcal{P}_{\mathcal{A}}$ such that the following holds: for any run $\rho = c_0 a_1 c_1 \dots a_n c_n$ of \mathcal{A} , there exists a run $\pi = x_0 a_1 x_1 \dots a_n x_n$ in $\mathcal{P}_{\mathcal{A}}$ such that for all $c_i = (G_i, \vec{d}_i)$ and $x_i = (s_i, w_i, \vec{d}_i)$ we have $\vec{d}_i = \vec{d}_i'$ and $G_i \triangleright w_i$ ($0 \leq i \leq n$), and vice versa.*

Proof. We assert that the semantics of $\mathcal{P}_{\mathcal{A}}$ is the usual semantics for pushdown systems with data, i.e., an infinite transition system with configurations $c = (y, w, d) \in Y \times S^* \times \mathbb{D}^{\mathcal{X}}$. Thus, we can interpret configurations also as follows: $(x, d) \in S^* \times \mathbb{D}^{\mathcal{X}}$ with $x = w \cdot y \in S^* \cdot (S \cup \{\varepsilon\})$.

Let $(G, \vec{d}) \in \text{Reach}(\mathcal{A})$ be reachable by a run $(G_0, \vec{d}_0) a_1 (G_1, \vec{d}_1) a_2 \dots a_n (G_n, \vec{d}_n)$. Then we can induce a run $(x_0, \vec{d}_0) a_1 (x_1, \vec{d}_1) a_2 \dots a_n (x_n, \vec{d}_n)$ in $\mathcal{P}_{\mathcal{A}}$ such that $\vec{d}_i = \vec{d}_i'$ and $G_i \triangleright x_i$ for $0 \leq i \leq n$.

By construction of $\mathcal{P}_{\mathcal{A}}$, $x_0 \triangleright G_0$ and $\vec{d}_0 = \vec{d}_0'$. We now assume that there exists a prefix of the QDAS's run of length $0 \leq j \leq n$ of the form $(G_0, \vec{d}_0) \dots (G_j, \vec{d}_j)$ such that there exists a run of the pushdown system $(x_0, \vec{d}_0) \dots (x_j, \vec{d}_j)$ that fullfills the induction hypothesis. We now consider the outcome of a QDAS transition labeled a_{j+1} . We know that G_j must be a path of vertices $v_0 \dots v_n$ connected by wait edges.

Sync. dispatch: dispatching a block γ on queue q leads to (G_{j+1}, \vec{d}_{j+1}) with $\vec{d}_j = \vec{d}_{j+1}$ and G_{j+1} is a path graph $v_0 v_1 \dots v_n v_{n+1}$ with new distinct vertex v_{n+1} where $\text{state}(v_{n+1}) = v_\gamma^0$. We mapped the dispatch rule to a **push** of the current state to the pushdown and jumping to the new initial state, i.e., we go from (x_j, \vec{d}_j) to (x_{j+1}, \vec{d}_{j+1}) where $\vec{d}_j = \vec{d}_{j+1}$ and $x_{j+1} = x_j \cdot s_\gamma^0$. Obviously, $G_{j+1} \triangleright x_{j+1}$.

Test/Assignment: G_{j+1} equals G_j except for $\text{state}_j(v_n) = s$ and $\text{state}_{j+1}(v_n) = s'$ and a possible change of \vec{d}_{j+1} according to the underlying data action.

Executing the same action on $\mathcal{P}_{\mathcal{A}}$ assures that $\vec{d}_{j+1} = \vec{d}_{j+1}$ and changing the control state of the pushdown only changes $x_j = w \cdot s$ to $x_{j+1} = w \cdot s'$; thus, $G_{j+1} \triangleright x_{j+1}$.

Termination: To apply the action G_j consists of a (non-empty) path ending in v with $\text{state}_j(v) \in F$ and $G_{j+1} = G_j \setminus v$, and $\vec{d}_j = \vec{d}_{j+1}$. Note that G_{j+1} could be possibly empty. Given a (x_j, \vec{d}_j) according to the induction hypothesis, then we have to consider two cases: either $x_j = w_j \cdot y_j$ with $w_j \in S^+$ and $y_j \in S$ (i.e., there is at least one element on the stack), or $x_j = y_j \in S$ (i.e., stack is empty). In the second case, we know that $x_j \in S_{\text{main}}$ and by the induction hypothesis, that $x_j = s_{\text{main}}^0$ and G_j a path of length 1. Now, $\mathcal{P}_{\mathcal{A}}$ takes the **empty?** transition leading to the (bottom) state ε , i.e., $x_{j+1} = \varepsilon$, hence G_{j+1} is empty and $G_{j+1} \triangleright \varepsilon$. If the stack is not empty, then we can take a **pop** transition such that $x_{j+1} = w \in S^+$ for $x_j = w \cdot s$, hence $G_{j+1} \triangleright x_{j+1}$. Obviously $\vec{d}_{j+1} = \vec{d}_j = \vec{d}_{j+1}$.

(Recall that we asserted dispatch and scheduling/dequeueing to be atomic, so we do not need to consider other actions of the scheduler.)

The reverse direction follows analogously as the previous inductive construction used necessary *sufficient* steps. \square

Proposition 3 *Given a synchronous QDAS \mathcal{A} , then we can construct a pushdown system with data $\mathcal{P}_{\mathcal{A}}$ such that the following holds: for any run $\rho = c_0 a_1 c_1 \dots a_n c_n$ of \mathcal{A} , there exists a run $\pi = x_0 a_1 x_1 \dots a_n x_n$ in $\mathcal{P}_{\mathcal{A}}$ such that for all $c_i = (G_i, \vec{d}_i)$ and $x_i = (s_i, w_i, \vec{d}_i)$ we have $\vec{d}_i = \vec{d}_i$ and $G_i \triangleright w_i$ ($0 \leq i \leq n$), and vice versa.*

Proof. We assert that the semantics of $\mathcal{P}_{\mathcal{A}}$ is the usual semantics for pushdown systems with data, i.e., an infinite transition system with configurations $c = (y, w, d) \in Y \times S^* \times \mathbb{D}^{\mathcal{X}}$. Thus, we can interpret configurations also as follows: $(x, d) \in S^* \times \mathbb{D}^{\mathcal{X}}$ with $x = w \cdot y \in S^* \cdot (S \cup \{\varepsilon\})$.

Let $(G, \vec{d}) \in \text{Reach}(\mathcal{A})$ be reachable by a run $(G_0, \vec{d}_0) a_1 (G_1, \vec{d}_1) a_2 \dots a_n (G_n, \vec{d}_n)$. Then we can induce a run $(x_0, \vec{d}_0) a_1 (x_1, \vec{d}_1) a_2 \dots a_n (x_n, \vec{d}_n)$ in $\mathcal{P}_{\mathcal{A}}$ such that $\vec{d}_i = \vec{d}_i$ and $G_i \triangleright x_i$ for $0 \leq i \leq n$.

By construction of $\mathcal{P}_{\mathcal{A}}$, $x_0 \triangleright G_0$ and $\vec{d}_0 = \vec{d}_0$. We now assume that there exists a prefix of the QDAS's run of length $0 \leq j \leq n$ of the form $(G_0, \vec{d}_0) \dots (G_j, \vec{d}_j)$ such that there exists a run of the pushdown system $(x_0, \vec{d}_0) \dots (x_j, \vec{d}_j)$ that fullfills the induction hypothesis. We now consider the outcome of a QDAS transition labeled a_{j+1} . We know that G_j must be a path of vertices $v_0 \dots v_n$ connected by wait edges.

Sync. dispatch: dispatching a block γ on queue q leads to (G_{j+1}, \vec{d}_{j+1}) with $\vec{d}_j = \vec{d}_{j+1}$ and G_{j+1} is a path graph $v_0 v_1 \dots v_n v_{n+1}$ with new distinct vertex v_{n+1} where $\text{state}(v_{n+1}) = v_\gamma^0$. We mapped the dispatch rule to a **push** of the current state to the pushdown and jumping to the new initial state, i.e., we go from (x_j, \vec{d}_j) to (x_{j+1}, \vec{d}_{j+1}) where $\vec{d}_j = \vec{d}_{j+1}$ and $x_{j+1} = x_j \cdot s_\gamma^0$. Obviously, $G_{j+1} \triangleright x_{j+1}$.

Test/Assignment: G_{j+1} equals G_j except for $\text{state}_j(v_n) = s$ and $\text{state}_{j+1}(v_n) = s'$ and a possible change of \vec{d}_{j+1} according to the underlying data action.

Executing the same action on $\mathcal{P}_{\mathcal{A}}$ assures that $\vec{d}_{j+1} = \vec{d}_{j+1}$ and changing the control state of the pushdown only changes $x_j = w \cdot s$ to $x_{j+1} = w \cdot s'$; thus, $G_{j+1} \triangleright x_{j+1}$.

Termination: To apply the action G_j consists of a (non-empty) path ending in v with $\text{state}_j(v) \in F$ and $G_{j+1} = G_j \setminus v$, and $\vec{d}_j = \vec{d}_{j+1}$. Note that G_{j+1} could be possibly empty. Given a (x_j, \vec{d}_j) according to the induction hypothesis, then we have to consider two cases: either $x_j = w_j \cdot y_j$ with $w_j \in S^+$ and $y_j \in S$ (i.e., there is at least one element on the stack), or $x_j = y_j \in S$ (i.e., stack is empty). In the second case, we know that $x_j \in$

S_{main} and by the induction hypothesis, that $x_j = s_{main}^0$ and G_j a path of length 1. Now, \mathcal{P}_A takes the **empty?** transition leading to the (bottom) state ε , i.e., $x_{j+1} = \varepsilon$, hence G_{j+1} is empty and $G_{j+1} \triangleright \varepsilon$. If the stack is not empty, then we can take a **pop** transition such that $x_{j+1} = w \in S^+$ for $x_j = w \cdot s$, hence $G_{j+1} \triangleright x_j$. Obviously $\vec{d}_{j+1} = \vec{d}_j = \widehat{\vec{d}}_j = \widehat{\vec{d}}_{j+1}$.

(Recall that we asserted dispatch and scheduling/dequeueing to be atomic, so we do not need to consider other actions of the scheduler.)

The reverse direction follows analogously as the previous inductive construction used necessary *sufficient* steps. \square

Lemma 1 *Given a finite set S and a function $f : S \rightarrow \mathbb{N}$, then there exists a finite automaton \mathcal{F}_f with alphabet S of size exponential in $|S|$ and polynomial in (in the binary encoding of) $\max_{s \in S} f(s)$ such that $\mathcal{L}(\mathcal{F}_f) = \{w \in S^* : |w|_s \geq f(s) \text{ for all } s \in S\}$.*

Proof. Given a set S and a function $f : S \rightarrow \mathbb{N}$. Let $k = \max_{s \in S} f(s)$ (which must exist as S is finite). Then \mathcal{F}_f is the finite automaton $\langle Q, S, q^0, \Delta, q^f \rangle$ with states $Q = S \times \{0 \dots k\}$ (interpreted as an S -indexed vector of values in $0 \dots k$), an action alphabet S , the initial state is q^0 where $q^0(s) = f(s)$, the final state is q^f where $q^f(s) = 0$. The transitions of \mathcal{F}_f are defined as follows: $(q, s, q') \in \Delta$ iff $q'(s) = q(s) - 1$ for $q(s) > 1$, else $q'(s) = q(s)$, and for all $t \in S \setminus \{s\}$ we have $q'(t) = q(t)$. Thus each transition labeled by an action s reduces the “counter” $q(s)$ by one until zero and once arrived at zero, the counter $q(s)$ remains zero for any further s action. Further, the control structure of \mathcal{F}_f is acyclic (except for the loops at q^f), thus each run can visit each state in $Q \setminus \{q^f\}$.

If $w = a_1 \dots a_n \in \mathcal{L}(\mathcal{A})$ then it was accepted by a run $q_0 a_1 q_1 \dots a_n q_n$ where $q_0 = q^0$ and $q_n = q^f$. Due to our construction of Δ , it holds for $w = a_1 \dots a_n$ that $|w|_s \geq q_0(s) = f(s)$ for all $s \in S$. If $w \notin \mathcal{L}(\mathcal{A})$ then there exists a run $q_0 a_1 q_1 \dots a_n q_n$ where $q_0 = q^0$ and for $q_n \neq q^f$ it holds that there exists at least one $s \in S$ such that $q_n(s) > 0$, each transition (q_{i-1}, a_i, q_i) assures that $q_{i-1}(s) \geq q_i(s)$, hence $|w|_s < f(s)$ for at least one $s \in S$. \square

Proposition 4 *Given a synchronous QDAS \mathcal{A} with states S and a function $f : S \rightarrow \mathbb{N}$, then one can generate a PDS $\mathcal{P}_{\mathcal{A},f}$ of size exponential in \mathcal{A} and a state s of $\mathcal{P}_{\mathcal{A},f}$, s.t. $\mathcal{P}_{\mathcal{A},f}$ reaches s iff f is Parikh coverable in \mathcal{A} .*

Proof. [Prop. 4] First, we construct the PDS with data \mathcal{P}_A and states S as mentioned before. Then, we translate the PDS with data to a bisimilar PDS without data $\widehat{\mathcal{P}}_A = \langle \widehat{Y}, \widehat{y}^0, \widehat{\Phi}, \widehat{\Sigma}, \widehat{\Delta} \rangle$ by encoding all possible valuations of variables into the PDS’s states by the standard product construction, i.e., $\widehat{Y} = S \times (\mathcal{X} \times \mathbb{D})$. Given $y \in \widehat{Y}$, let $S(y) \in S$ denote the original state component. Note: $\widehat{\mathcal{P}}_A$ is at most exponentially larger as \mathcal{P}_A and this construction does not change the pushdown system’s behaviour with respect to the stack but only internal actions.

Second, from the function f , we construct the automaton $\mathcal{F}_f = \langle Q, S, q^0, \Delta_{\mathcal{F}}, q^f \rangle$ analogous to Lemma 1.

Finally, we define the PDS $\mathcal{P}_{\mathcal{A},f} = \langle Y, y^0, \Phi, \Sigma, \Delta_{\mathcal{A},f} \rangle$ as follows

- states are $Y = \widehat{Y} \cup Q$ (assuring disjointness by relabeling when necessary)
- $y^0 = \widehat{y}^0$ is the initial state
- $\Phi = \widehat{\Phi}$ is the stack alphabet (where $\widehat{\Phi} = S$ due to the above construction)
- $\Sigma = \widehat{\Sigma} \cup \{\varepsilon\}$
- a tuple (y, a, y') is a rule in $\Delta_{\mathcal{A},f} \subseteq Y \times \Sigma \times Y$ iff one of the following holds
 - $(y, a, y') \in \widehat{\Delta}$ (include all transition rules of $\widehat{\mathcal{P}}_{\mathcal{A}}$);
 - $a = \text{pop}(s)$ for $s \in \Phi$ and $(q, s, q') \in \Delta_{\mathcal{F}}$ (include rules of \mathcal{F}_f and change an s action to $\text{pop}(s)$ for $s \in S$);
 - $y \in \widehat{Y}$, $a = \text{push}(z)$ for $z = S(y)$, and $y' = q^0$ (connect all states in \widehat{Y} with the initial state of \mathcal{F}_f , additionally stocking the current “state”-component on the stack).

Note that $\mathcal{P}_{\mathcal{A},f}$ is of size exponential with respect to both the QDAS and f due to serial composition.

We now have to show that if there is a run in $\mathcal{P}_{\mathcal{A},f}$ that reaches the state q^f , then there exists configuration $c = (G, \vec{d})$ of \mathcal{A} such that $f \preceq \text{Parikh}(G)$.

Assert that there exists a run of $\mathcal{P}_{\mathcal{A},f}$ reaching q^f , then it must be of the following form $\langle x_0, a_1, x_1, \dots, a_k, x_k, a_{k+1}, x_{k+1}, a_{k+2}, \dots, a_n, x_n \rangle$ where $x_i = (y_i, w_i) \subseteq Y \times S^*$ are the corresponding infinite transition systems configurations. Further, $y_0 = y^0$, $y_n = q^f$, $y_{k+1} = q^0$, and $\langle y_1 \dots y_k \rangle$ is a subrun that only uses states in \widehat{Y} as well as transitions in $\widehat{\mathcal{P}}_{\mathcal{A}}$; $\{y_{k+1}, \dots, y_n\} \subseteq Q$ and the corresponding transitions are derived from $\Delta_{\mathcal{F}}$, as well as $a_{k+1} = \text{push}(S(y_k))$.

Let us take a closer look on the first part of the run: $\langle y_0, a_1, \dots, a_n, x_k \rangle$ is equivalent to a run of $\widehat{\mathcal{P}}_{\mathcal{A}}$ that reaches a configuration x_k . The latter is, following Propositions 3 and ??, similar to a run of the original QDAS \mathcal{A} that reaches a configuration $c = (G, \vec{d})$ where $G \triangleright y_k \cdot S(y_k)$. Thus, $c \in \text{Reach}(\mathcal{A})$.

The transition $(x_k, \text{push}(S(y_k)), x_{k+1})$ now transfers the encoding of G to the stack, i.e., $w_{k+1} = y_k \cdot S(y_k)$. All other information on data encoded in y_k is lost in this step.

Now, by Lemma 1 we know that the subrun $\langle x_{k+1}, a_{k+2}, \dots, a_n, x_n \rangle$ leading to the final state of \mathcal{F}_f assures that $|w_{k+1}|_s \geq f(s)$ for all $s \in S$. Hence, for the previously found $c = (G, \vec{d}) \in \text{Reach}(\mathcal{A})$ it holds that $f \preceq \text{Parikh}(G)$. \square

Let us take a closer look on the dispatches that happen in runs of synchronous QDAS that have only *serial* queues. Assume a run of such a QDAS, and suppose the first dispatch performed along this run (by **main**) is $\text{dispatch}_{\mathbf{s}}(q, \gamma)$. As the dispatch is synchronous, **main** is blocked, and the scheduler has to dequeue γ to let the system progress. Cleraly, if γ performs a synchronous dispatch $\text{dispatch}_{\mathbf{s}}(q, \gamma')$ to the same queue q , we reach a deadlock. Indeed, the task running γ is blocked by the synchronous dispatch of γ' , but we need to wait for the termination of γ to be able to dequeue γ' from q (because q is serial). So, γ has to dispatch its blocks to other queues. For the same reason, we also reach

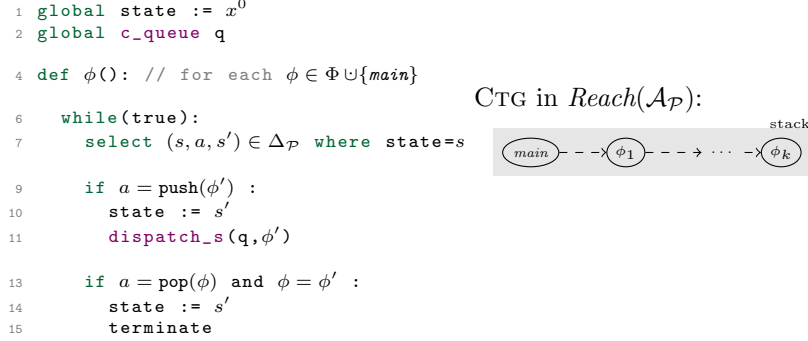


Figure 5: From a pushdown system to a QDAS: $main$ and ϕ for $\phi \in \Phi$ a deadlock if a block called by γ performs a synchronous dispatch into q . We conclude that, in all reachable CTG, the following holds for all queues: either the queue contains one block and there is no running task from this queue, or the queue is empty, and there is at most one running task from this queue. Hence, all the reachable CTG have at most $|SQID| + 2$ vertices. Thus, the pushdown systems used in all previous constructions have bounded stack height and we can apply the emptiness test on a finite state system when proving Proposition 4. The lower bound can be derived from Proposition 2. Thus we can derive:

Proposition 4 *Given a synchronous QDAS \mathcal{A} with states S and a function $f : S \rightarrow \mathbb{N}$, then one can generate a PDS $\mathcal{P}_{\mathcal{A},f}$ of size exponential in \mathcal{A} and a state s of $\mathcal{P}_{\mathcal{A},f}$, s.t. $\mathcal{P}_{\mathcal{A},f}$ reaches s iff f is Parikh coverable in \mathcal{A} .*

B.0.1 From Pds to Qdas

Given a PDS \mathcal{P} , we construct a synchronous QDAS $\mathcal{A}_{\mathcal{P}}$ as shown in Figure 5. The underlying idea is the inverse of the above simulation: we map a **push** action of a letter ϕ to synchronous dispatch call of a block ϕ and simulate the stack contents in the CTG such that we can only map a **pop** action to a task's termination if we match the topmost letter of the stack, encoded in the block name.

The control state of the PDS is stored in the variable **state** and the behaviour of the control structure of \mathcal{P} is encoded as non-deterministic choice (line 7) that assures that reaching the dispatch and termination actions (lines 11/15) demands that the selected transition rule harmonizes with the current change of the variable **state** from s to s' and that a **push**(ϕ') action is only possible if the currently running task is labeled by the blockname ϕ' (line 13).

A reachable configuration of $\mathcal{A}_{\mathcal{P}}$ is given by (G, \vec{d}) where G is—as discussed before—a path of vertices $v_0 v_1 \dots v_k$. As before, synchronous dispatch calls assure there is no more than one task active at the same time. Given $c = (G, \vec{d}) \in \text{Reach}(\mathcal{A}_{\mathcal{P}})$ and a configuration $y = (x, w) \in X \times \Phi^*$ that is reachable in \mathcal{P} ; then c is represented by y , written $c \triangleright y$, iff $\vec{d}(\text{state}) = x$ and for $w = w_1 \dots w_k$ $\lambda(v_i) = w_i$ for $1 \leq i \leq k$ and $\lambda(v_0) = \text{main}$. Hence, the state of the PDS is stored in the variable **state**, and the path $v_1 \dots v_k$ encodes in the underlying

task's blocks the stack content, where the empty stack is represented by a single vertex labeled by *main*.

Proposition 9 *Given a pushdown system \mathcal{P} , then we can generate a synchronous QDAS $\mathcal{A}_{\mathcal{P}}$ such that the following holds: for any run $\pi = y_0 a_1 y_1 \dots a_n y_n$ in \mathcal{P} there exists a run $\rho = c_0 a_1 c_1 \dots a_n c_n$ of $\mathcal{A}_{\mathcal{P}}$ such that for all $c_i \triangleright x_i$ ($0 \leq i \leq n$), and vice versa.*

Proof. Given a run $\rho = y_0 a_1 y_1 \dots a_k y_k$ of the PDS \mathcal{P} . W.l.o.g. let us consider in the following underlying sequence of configurations and fired transition rules $y_0 \delta_1 y_1 \dots \delta_k y_k$ where $\delta_i = (x_i, a_i, x'_i) \in \Delta_{\mathcal{P}}$ for $1 \leq i \leq k$.

We show inductively how $\mathcal{A}_{\mathcal{P}}$ generates a run that simulates ρ .

For the initial configuration of \mathcal{P} $y_0 = (x^0, \varepsilon)$ and the initial configuration $c_0 = (G, \vec{d})$ with G consists of a single node v_0 with $\lambda(v_0) = \text{main}$ and $\vec{d}(\text{state}) = x^0$ it holds that $c_0 \triangleright y_0$.

Now assert that the PDS \mathcal{P} reached configuration y_i ($0 \leq i \leq k$) such that $\mathcal{A}_{\mathcal{P}}$ simulated the prefix of the run until $c_i = (G_i, \vec{d}_i)$ with $c_i \triangleright y_i$. Assert that G_i is a path $v_0 v_1 \dots v_l$. We do a case-by-case analysis with respect to $\delta_{i+1} = (x, a, x')$ that leads to y_{i+1} :

- only the task corresponding to v_l is active and the only way to exit its **while** loop is via the lines 11 and 15, that assure that line 7 selected $\delta = (x, a, x') \in \Delta_{\mathcal{P}}$ with $\vec{d}_i(\text{state}) = x$, and that we set $\vec{d}_{i+1}(\text{state}) = x'$;
- if $a = \text{push}(\phi)$ for $\phi \in \Phi$, then we fire the synchronous dispatch that leads to $G_{i+1} = v_0 \dots v_l v_{l+1}$ with $\lambda(v_{l+1}) = \phi$, thus $(G_{i+1}, \vec{d}_{i+1}) \triangleright y_{i+1}$;
- if $a = \text{pop}(\phi)$ for $\phi \in \Phi$ and we left the while loop then $\lambda(v_l) = \phi$ (by line 13), and G_{i+1} equals $v_0 \dots v_{l-1}$, thus $(G_{i+1}, \vec{d}_{i+1}) \triangleright y_{i+1}$.

The reverse direction follows analogously by considering lines 10; 11 and 14; 15 as atomic actions (i.e., setting the **state** variable and changing the call graph of the QDAS).

B.1 Asynchronous Concurrent Qdas

Proposition 5 *For all concurrent asynchronous QDAS \mathcal{A} with set of location S , we can build, in polynomial time, a Petri net $N_{\mathcal{A}}$ s.t. f is Parikh-coverable in \mathcal{A} iff $m \in \text{Cover}(N_{\mathcal{A}})$, where m is the marking s.t. for all $s \in S$: $m(s) = f(s)$ and for all $p \in P \setminus S$: $m(p) = 0$.*

The proof of the proposition relies on the following lemma, showing that $N_{\mathcal{A}}$ can simulate precisely the sequence of Parikh images that are reachable in \mathcal{A} . Let (G, \vec{d}) be a configuration of \mathcal{A} , and let m be marking of $N_{\mathcal{A}}$. We say that m *encodes* (G, \vec{d}) , written $m \triangleright (G, \vec{d})$ iff: (i) for all $x \in \mathcal{X}$: $m(x, \vec{d}(x)) = 1$, (ii) for all $x \in \mathcal{X}$: for all $d \in \mathbb{D} \setminus \{\vec{d}(x)\}$: $m(x, d) = 0$ and (iii) for all $s \in S$ $m(s) = \text{Parikh}(G)(s)$. Then:

Lemma 2 *Let \mathcal{A} be a concurrent asynchronous QDAS with set of variables \mathcal{X} and set of locations S , and let $N_{\mathcal{A}}$ be its associated PN. Then, for all $(G, \vec{d}) \in \text{Reach}(\mathcal{A})$ there is $m \in \text{Reach}(N_{\mathcal{A}})$ s.t. $m \triangleright (G, \vec{d})$ and for all $m \in \text{Reach}(N_{\mathcal{A}})$, there is $(G, \vec{d}) \in \text{Reach}(\mathcal{A})$ s.t. $m \triangleright (G, \vec{d})$.*

Proof. We prove the two statements separately.

Let (G, \vec{d}) be a configuration in $\text{Reach}(\mathcal{A}_N)$, and let $(G_0, \vec{d}_0)a_0(G_1, \vec{d}_1)a_1 \cdots a_{n-1}(G_n, \vec{d}_n)$ be a run s.t. $(G, \vec{d}) = (G_n, \vec{d}_n)$. Let us build, inductively, a run $m_0m_1 \cdots m_k$ of $N_{\mathcal{A}}$ s.t. $m_k \triangleright (G, \vec{d})$. The induction is on the length n of the QDAS run.

Base case $n = 0$. It is easy to check that $m_0 \triangleright (G_0, \vec{d}_0)$.

Inductive case $n = \ell$. Let us assume that $m_0m_1 \cdots m_j$ is a run of $N_{\mathcal{A}}$ s.t. $m_j \triangleright (G_{\ell-1}, \vec{d}_{\ell-1})$, and let us show how to complete it, if needed. We consider several case depending on a_{n-1} . In the case where $a_{n-1} = \varepsilon$ and the scheduler action consists in dequeuing a block from a queue, we have $\text{Parikh}(G_{\ell-1}) = \text{Parikh}(G_{\ell})$ and $\vec{d}_{\ell} = \vec{d}_{\ell-1}$. By induction hypothesis $m_j \triangleright (G_{\ell-1}, \vec{d}_{\ell-1})$, hence $m_j \triangleright (G_{\ell}, \vec{d}_{\ell})$, and we do not add elements to the run built so far. In the case where $a_{\ell-1} = \text{dispatch}_{\mathbf{a}}(\gamma, q)$, we assume $(s, a_{\ell-1}, s') \in \Delta$ is the corresponding LTS transition. Clearly, $\text{Parikh}(G_{\ell})(s') = \text{Parikh}(G_{\ell-1})(s') + 1$, $\text{Parikh}(G_{\ell})(s) = \text{Parikh}(G_{\ell-1})(s) - 1$, $\text{Parikh}(G_{\ell})(s_{\gamma}^0) = \text{Parikh}(G_{\ell-1})(s_{\gamma}^0) + 1$ and for all other location s : $\text{Parikh}(G_{\ell})(s) = \text{Parikh}(G_{\ell-1})(s)$. It is easy to check that the PN transition t s.t. $I(t)(p) = 1$ iff $p = s$ and $O(t)(p) = 1$ iff $p \in \{s', s_{\gamma}^0\}$ is fireable from m_j (as $m_j \triangleright (G_{\ell-1}, \vec{d}_{\ell-1})$ by induction hypothesis) and yields the same effect, i.e. the marking m with $m_j \xrightarrow{t} m$ is s.t. $m \triangleright (G_{\ell}, \vec{d}_{\ell})$. All the other cases (test, assignment and task termination) are treated similarly.

Now, let $m_0m_1 \cdots m_n$ be a run of $N_{\mathcal{A}}$ and let us build, inductively, a run $(G_0, \vec{d}_0)a_0(G_1, \vec{d}_1)a_1 \cdots a_{k-1}(G_k, \vec{d}_k)$ s.t. $m_n \triangleright (G_k, \vec{d}_k)$ and all the queues are empty in G_n . The induction is on the length n of the PN run.

Base case $n = 0$. It is easy to check that $m_0 \triangleright (G_0, \vec{d}_0)$.

Inductive case $n = \ell$. Let us assume that $(G_0, \vec{d}_0)a_0 \cdots a_{j-1}(G_j, \vec{d}_j)$ is a run of \mathcal{A} s.t. $m_{\ell-1} \triangleright (G_j, \vec{d}_j)$ and all the queues are empty in G_j . Let t be the PN transition s.t. $m_{\ell-1} \xrightarrow{t} m_{\ell}$ and let us show how we can extend the run of \mathcal{A} . We consider several cases. If t is a transition that corresponds to an asynchronous dispatch, then there are s, s', γ and q s.t. $I_t(p) = 1$ iff $p = s$ and $O_t(p) = 1$ iff $p \in \{s', s_{\gamma}^0\}$. By definition of $N_{\mathcal{A}}$, there is a transition $(s, \text{dispatch}_{\mathbf{a}}(\gamma, q), s')$ in \mathcal{A} . Moreover, $m_{\ell-1}(s) \geq 1$, since t is fireable from $m_{\ell-1}$. As $m_{\ell-1} \triangleright (G_j, \vec{d}_j)$, the $(s, \text{dispatch}_{\mathbf{a}}(\gamma, q), s')$ is fireable from (G_j, \vec{d}_j) , and leads to a configuration (G_{j+1}, \vec{d}_{j+1}) , where a γ block has been enqueued in q , hence $\vec{d}_{j+1} = \vec{d}_j$, $\text{Parikh}(G_{j+1})(s) = \text{Parikh}(G_j)(s) - 1$, $\text{Parikh}(G_{j+1})(s') = \text{Parikh}(G_j)(s') + 1$, $\text{Parikh}(G_{j+1})(s_{\gamma}^0) = \text{Parikh}(G_j)(s_{\gamma}^0) + 1$ and for all other state s'' : $\text{Parikh}(G_{j+1})(s'') = \text{Parikh}(G_j)(s'')$. It is easy to check that $m_{\ell} \triangleright (G_{j+1}, \vec{d}_{j+1})$, however, queue q contains a call to γ in G_{j+1} and is thus the only non-empty queue in this CTG. Thus, from (G_{j+1}, \vec{d}_{j+1}) , we execute the scheduler action that dequeues from q . This has no effect on the Parikh image




Figure 6: The LTS of bloc p .

of the CTG. Thus, we reach (G_{j+2}, \vec{d}_{j+2}) s.t. $\vec{d}_{j+1} = \vec{d}_{j+2}$, $\text{Parikh}(G_{j+1}) = \text{Parikh}(G_{j+2})$, hence $m \ell \triangleright (G_{j+2}, \vec{d}_{j+2})$ too, and all the queues are empty in G_{j+2} , which concludes the induction step. All the other cases are treated similarly. \square

We can now prove Proposition 5: *Proof.* It is easy to check that the construction of $N_{\mathcal{A}}$, as described above, is polynomial. Then, assume f is Parikh coverable in \mathcal{A} , i.e. there is $(G, \vec{d}) \in \text{Reach}(\mathcal{A})$ s.t. $f \preceq \text{Parikh}(G)$. By Lemma 2, there is $m' \in \text{Reach}(N_{\mathcal{A}})$ s.t. $m' \triangleright (G, \vec{d})$. Hence, for all $s \in S$: $m'(s) = \text{Parikh}(G)(s)$. So, for all $s \in S$: $m(s) = f(s) \leq \text{Parikh}(G)(s) = m'(s)$. Hence, $m \preceq m'$ (as $m(p) = 0$ for all $p \notin S$). Since $m' \in \text{Reach}(N_{\mathcal{A}})$, we conclude that $m \in \text{Cover}(N_{\mathcal{A}})$. On the other hand, assume $m \in \text{Cover}(N_{\mathcal{A}})$, with $m(p) = 0$ for all $p \notin S$, and let f be s.t. for all $s \in S$: $f(s) = m(s)$. Since $m \in \text{Cover}(N_{\mathcal{A}})$, there is $m' \in \text{Reach}(N_{\mathcal{A}})$ s.t. $m \preceq m'$. By Lemma 2, there is $(G, \vec{d}) \in \text{Reach}(\mathcal{A})$ s.t. $m' \triangleright (G, \vec{d})$. Thus, by definition of \triangleright , for all $s \in S$: $m'(s) = \text{Parikh}(G)(s)$. Thus, since $m \preceq m'$ and by definition of f , we conclude that for all $s \in S$: $f(s) = m(s) \leq m'(s) = \text{Parikh}(G)(s)$. Hence, f is Parikh-coverable in \mathcal{A} . \square

Proposition 6 *For all Petri nets N , we can build, in polynomial time, a concurrent asynchronous QDAS \mathcal{A}_N s.t. $m \in \text{Cover}(N)$ iff there exists $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_N)$ with $G \triangleright m$.*

The proof of Proposition 6 is split into two lemmata, given hereunder. They rely on an alternate characterization of $\text{Cover}(N)$. That is, $m \in \text{Cover}(N)$ iff m is reachable by a so-called *lossy* run of N , i.e. a sequence of markings $m'_0 m'_1 \dots m'_n$ s.t. $m'_0 \preceq m_0$ and for all $0 \leq i \leq n-1$: there is \bar{m}_{i+1} and a transition t_i s.t. $m'_i \xrightarrow{t_i} \bar{m}_{i+1}$ and $m'_{i+1} \preceq \bar{m}_{i+1}$. Intuitively, a lossy run corresponds to firing a transition of the PN, and then spontaneously losing some tokens. The proof of these lemmata also assumes that each $p \in P$, the LTS $\mathcal{TS}_p = \langle \{s_p^0, s_p^{mid}, s_p^{fin}\}, s_p^0, \Sigma, \Rightarrow \rangle$ is as depicted in Fig. 6.

Lemma 3 *Let $N = \langle P, T, m_0 \rangle$ be a PN, and let $\mathcal{A}_N = \langle CQID, \emptyset, \Gamma, \text{main}, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$ be its corresponding QDAS. **If** $m \in \text{Cover}(N)$ **then** there exists $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_N)$ s.t. $G \triangleright m$.*

Proof. Let m be a marking from $\text{Cover}(N)$. and let $m'_0 m'_1 \dots m'_n$ be a lossy PN run s.t. $m = m_n$. The proof is by induction on the length of the run. More precisely, we show that, for all $0 \leq i \leq n$, there is a reachable configuration $(G_i, \vec{d}_i) \in \text{Reach}(\mathcal{A}_N)$ s.t.: for all $p \in P$: $\vec{d}_i(v_p) = 0$, $G_i = \langle V^i, E^i, \lambda^i, \text{queue}^i, \text{state}^i \rangle$, $G_i \triangleright m$, $m \preceq m'_i$ and $E^i = \emptyset$.

Base case: m'_0 . Let us consider the run of \mathcal{A}_N that consists in: (a) executing block **main** up to line 8, then (b) emptying the queue C . The execution of (a) has the effect that: (i) all v_p variables are initialized to 0 and keep this value, (ii) for all place p : at most $m_0(p)$ copies of block **p** are asynchronously dispatched in queue C and (iii) one copy of block **trans** is dispatched in C . Then, the execution of (b) creates one running task for each block that is present in C . Thus, the execution of (a) followed by (b) reaches a configuration (G_0, \vec{d}_0) with $G_0 = \langle V^0 = V_T^0 \cup V_C^0, E^0, \lambda^0, queue^0, state^0 \rangle$ s.t. $V_C^0 = \emptyset$ (the queue has been emptied), for all p : $|\{v \in V_T^0 \mid \lambda(v) = \mathbf{p}\}| = m_0(p)$, $|\{v \in V_T^0 \mid \lambda(v) = \mathbf{trans}\}| = 1$ and $E^0 = \emptyset$ (the queue is empty and all the calls are asynchronous). Moreover, $state$ is such that each task running a **p** block is still in its initial state s_p^0 , hence $G_0 \triangleright m_0$. Similarly, the task running the **trans**() block is about to enter the **while** loop at line 14. Finally, as the variables have been initialized to 0 and not modified, we have $\vec{d}_0(v_p) = 0$ for all $p \in P$.

Inductive case: m_i Let us assume there exist $(G_{i-1}, \vec{d}_{i-1}) \in Reach(\mathcal{A}_N)$ that respects all the conditions given at the beginning of the proof (in particular $G_{i-1} \triangleright m_{i-1}$). Let t_i and \bar{m}_i be the PN transition and marking s.t. $m_{i-1} \xrightarrow{t_i} \bar{m}_i$ and $m_i \preceq \bar{m}_i$ and let us show that \mathcal{A}_N can simulate it. This is achieved by the following sequence of actions in \mathcal{A}_N . First, the block executing **trans** enters the **while** loop at line 14 and selects t_i as transition t . Then, it sets all the variables v_p s.t. $I_{t_i}(p) = 1$ to 1. Thus, at that point v_p contains 1 iff $I_{t_i}(p) = 1$, since all v_p variables were equal to 0 by induction hypothesis. Then, the task executing **trans** is blocked as it need to wait up to the point where all v_p are equal to 0. Since $G_{i-1} \triangleright m_{i-1}$ by induction hypothesis, we know that there are, in G_{i-1} , $m_{i-1}(p)$ tasks executing block **p**, for all $p \in P$. However, t_i is fireable from m_{i-1} , and a loss of $\bar{m}_i - m_i$ token is still possible after the firing. Hence, $m_{i-1}(p) \geq (I_{t_i}(p) + \bar{m}_i(p) - m_i(p))$ for all p . Thus, for all p , there is at least $(I_{t_i}(p) + \bar{m}_i(p) - m_i(p))$ tasks executing **p** in G_{i-1} . Thus, we complete the run of \mathcal{A}_N by letting, for all p , $(I_{t_i}(p) + \bar{m}_i(p) - m_i(p))$ **p** task execute lines 11 in turn one after the other. Then, letting them all execute line 12, and reach their final state (Remark that all the **p** task must first execute line 11 before one of them can execute line 12, as this sets v_p to 0 and would prevent other tasks to execute line 11). This is possible because none of those tasks are blocked, since the CTG contains no edge, by induction hypothesis. At that point, \mathcal{A}_N has reached a configuration (G', \vec{d}') s.t. $\vec{d}'(v_p) = 0$ for all $p \in P$ (by line 12) and where $G' \triangleright m_{i-1} - (I_{t_i} + \bar{m}_i - m_i)$. Moreover, G' still respects all the other hypothesis as no new dispatch have been performed. Then, the simulation of t_i proceeds by letting the **trans** task finish the current iteration of the main **while** loop. This consists in executing the **for** loop of line 19, which dispatches one **p** block in C iff $O_{t_i}(p) = 1$, i.e., the effect of t_i is to add a token to p . Finally, the scheduler empties queue C and creates tasks for all the blocks that have just been added to C . It also kills all the **p** tasks that have reached their final state. As a consequence, the configuration that is reached is (G_i, \vec{d}_i) , where $G_i \triangleright m_{i-1} - (I_{t_i} + \bar{m}_i - m_i) + O_{t_i} = (m_{i-1} - I_{t_i} + O_{t_i}) - \bar{m}_i + m_i = \bar{m}_i - \bar{m}_i + m_i = m_i$

and \vec{d}_i is s.t. $\vec{d}_i(v_p) = 0$ for all $p \in P$. Moreover, since the queue has been emptied by the scheduler, G_i contains only task nodes and no edge, as all the calls are asynchronous. The task executing **trans** is still active and at line 14, and all the **p** tasks are in their initial state. \square

Lemma 4 *Let $N = \langle P, T, m_0 \rangle$ be a PN, and let $\mathcal{A}_N = \langle CQID, \emptyset, \Gamma, \text{main}, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$ be its corresponding QDAS. **If** there are $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_N)$ and m s.t. $G \triangleright m$ **then** $m(G) \in \text{Cover}(N)$.*

Proof. For a CTG G of \mathcal{A}_N with set of vertices V , we denote by $M(G)$ the marking of N s.t. for all $p \in P$: $M(G)(p) = |\{v \in V \mid \text{state}(v) = s_p^0\}|$. Thus, in the case where G encodes a configuration s.t. **trans** is at line 14, **main** is at line 8, and all the **p** blocks are in their initial state, then $G \triangleright M(G)$.

In order to establish the lemma, we prove a stronger statement: every time we reach, along a run, a configuration (G, \vec{d}) s.t. **trans** is at line 14, then $M(G) \in \text{Cover}(N)$. Formally, let $\rho = (G_0, \vec{d}_0) a_0 (G_1, \vec{d}_1) a_1 (G_2, \vec{d}_2) \cdots (G_n, \vec{d}_n)$ be a run of \mathcal{A}_N , where, for all $0 \leq i \leq n$: $G_i = \langle V_i, E_i, \lambda_i, \text{queue}_i, \text{state}_i \rangle$. Let $\pi : \{0, \dots, k\} \rightarrow \{0, \dots, n\}$ be the monotonically increasing function s.t. $k \leq n$ and for all $0 \leq j \leq n$: there exists $v \in V_j$ with $\text{state}_i(v) = s_{\text{trans}}^{14}$ iff there is $0 \leq \ell \leq k$ with $k = \pi(\ell)$. That is the sequence $\pi(1), \pi(2), \dots, \pi(k)$ identifies the indexes of all the configurations of the run where **trans** is at line 14. Let us show, by induction on i that all the $M(G_{\pi(i)})$'s are reachable in the lossy semantics of N .

Base case $i = 0$ Let us show that $M(G_{\pi(0)}) = m_0$, i.e., that the first time **trans** reaches line 14, $M(G_{\pi(0)})$ is the initial marking of N . Observe that the prefix of the run must have the following form. Initially, only the **main** block is executing: it first sets all the variables v_p to 0, then dispatches asynchronously at most $m_0(p)$ calls to each **p** block (for all $p \in P$), then finally dispatches an asynchronous call to **trans** and reaches line 8. Along this execution, the scheduler might decide to pick up some **p** blocks from C . However, as long as the scheduler has not scheduled the call to **trans**, the CTG met along the run do not encode any marking, by definition of \triangleright . When the scheduler starts a task to run the **trans** block, we thus reach a configuration (G, \vec{d}) where: (i) the queue C is empty, as dequeuing the **trans** block is possible only if all the **p** blocks have been dequeued, and no other dispatch has been performed; (ii) all the **p** tasks are blocked in their initial state as $\vec{d}(v_p) = 0$ for all $p \in P$; and (iii) **main** is still blocked in the infinite loop at line 8. Since the scheduler has just dequeued **trans** from C , G is necessarily the first CTG to encode a marking, so $G = G_{\pi(0)}$. Moreover, by the loop at line 4, it is clear that $G \triangleright m$ with $m \preceq m_0$.

Inductive case $i = \ell \geq 1$ The induction hypothesis is that $M(G_{\pi(i-1)}) \in \text{Cover}(N)$. Let us consider the $\rho' = (G_{\pi(\ell-1)}, \vec{d}_{\pi(\ell-1)}) \cdots (G_{\pi(\ell)}, \vec{d}_{\pi(\ell)})$, i.e. the portion of ρ that allows to reach $(G_{\pi(\ell)}, \vec{d}_{\pi(\ell)})$ from $(G_{\pi(\ell-1)}, \vec{d}_{\pi(\ell-1)})$. We consider two cases:

1. Either **trans** has not performed an iteration of its main **while** loop along ρ' . In this case, the only actions that can occur along ρ' are scheduler

actions consisting in dequeuing **p** blocks or the termination of some **p** tasks that were still in state s_p^{mid} . In both cases, this does not modify the value of $M(G)$, so $M(G_{\pi(i)}) = M(G_{\pi(i-1)}) \in \text{Cover}(N)$.

2. Or **trans** has performed a complete iteration of its main **while** loop possibly interleaved with the dequeue of **p** blocks and the termination of **p** tasks. Since the dequeues and terminations have no influence on the value of $M(G)$ as argued above, let us focus on the effect of executing one iteration of the **while** loop. The iteration first selects a PN transition t and sets all the variables v_p s.t. $I_t(p) = 1$ to 0. The reached configuration is then (G, \vec{d}) where $M(G) = M(G_{\pi(i-1)})$, as these operations do not manipulate **p** blocks or tasks. Then, **trans** is blocked by the test at line 18. As only **p** blocks can set v_p variables to 0, we are sure that, when **trans** reaches line 19, *at least* $I_t(p)$ **p** blocks have left their initial state, for all $p \in P$. Thus, when **trans** is at line 19, the configuration is (G', \vec{d}') , where for all $p \in P$: $M(G')(p) \leq M(G)(p) - I_t(p) = M(G_{\pi(i-1)}) - I_t(p)$. Afterwards, **trans** terminates the iteration of the **while** loop by dispatching $O_t(p)$ **p** blocks for all $p \in P$, and reaches line 14, which finishes ρ' . Hence, we reach $(G_{\pi(i)}, \vec{d}_{\pi(i)})$, where for all $p \in P$: $M(G_{\pi(i)})(p) \leq M(G_{\pi(i-1)}) - I_t(p) + O_t(p)$. Since $M(G_{\pi(i-1)}) \in \text{Cover}(N)$ by induction hypothesis, we conclude that $M(G_{\pi(i)}) \in \text{Cover}(N)$ too. \square

B.2 Asynchronous Serial Qdas

We establish the undecidability for asynchronous serial QDAS by a reduction from the control-state reachability problem in a fifo system. Let $F = \langle S_F, s_F^0, M, \Delta_F \rangle$ be a fifo system and let $c \in S_F$ be a control state whose reachability has to be tested. We build the asynchronous serial QDAS $\mathcal{A}_F = \langle \emptyset, \{q\}, \Gamma, \text{main}, \mathcal{X}, \Sigma, (\mathcal{TS}_\gamma)_{\gamma \in \Gamma} \rangle$ on domain $\mathbb{D} = M \cup S_F \cup \{\varepsilon\}$, where $\Gamma = M \cup \{\varepsilon, \text{main}\}$, $\mathcal{X} = \{\text{state}, \text{head}\}$ and the \mathcal{TS}_γ are given by the pseudo code in Fig. 7.

Intuitively, runs of \mathcal{A}_F simulate the runs of F , by encoding the current state of F in variable **state** and the content of F 's queue into the content of the serial queue **q**. More precisely, it is easy to check that, once **main** has reached line 8, all the CTG that are reached in \mathcal{A}_F are of either shapes depicted in Fig. 7, for $\{m_1, \dots, m_n, m\} \subseteq M \cup \{\varepsilon\}$. That is, there are at most two running tasks: **main** and possibly one task running a m block (for $m \in M \cup \{\varepsilon\}$), that has to terminate to allow a further dequeue from **q**. This is because **q** is a serial queue and all the dispatches are asynchronous. When the CTG is of shape (b), the duty of the running m block is to simulate a run of F . It runs an infinite **while** loop (line 11 onwards – ignore the test at line 10 for the moment), that (i) tests whether c has been reached (line 12) and jumps to line 20 if it is the case; (ii) guesses a transition (s, a, s') of F ; and (iii) checks that the guessed transition is indeed fireable from the current configuration of F , and, if yes, simulate it. This consists in, first testing that s is the current state (line 14). If not, the block jumps to the infinite loop of line 19, which ends the simulation. Otherwise, the current state is update to s' , and the channel operation is then simulated. A

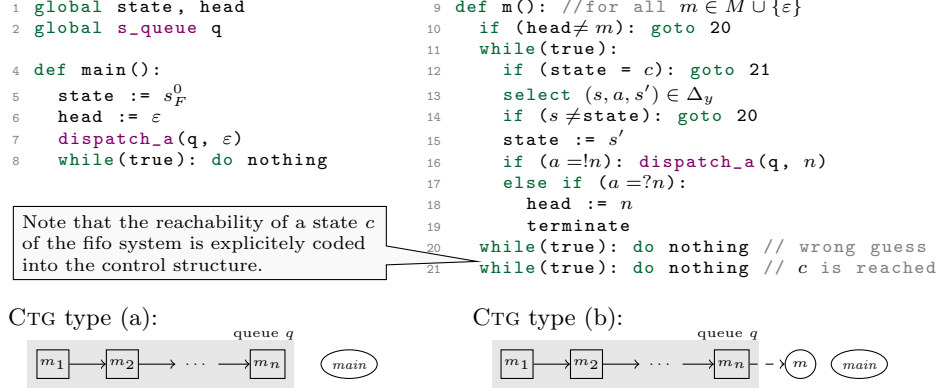


Figure 7: Fifo system encoding into a serial asynchronous QDAS/ two types of CTG in this case

send of message m is simulated (line 16) by an asynchronous dispatch of block m to q . The simulation of a receive of m from q is more involved, as only the scheduler can decide to dequeue a block from q , and this can happen only if the current running block terminates (line 19). Still, we have to check that message m is indeed in the head of q . This is achieved by setting global variable **head** to m , and letting the next dequeues block check that itself encodes the value stored into **head**. This is performed at line 10. If this test is not satisfied, the block jumps to the infinite loop of line 20, and the simulation ends. Otherwise, it proceeds with the simulation. Thus, in all reachable configurations of \mathcal{A}_F , a block m (with $m \in M \cup \{\varepsilon\}$) will reach line 21 iff c is reachable in F . This effectively reduces the control location reachability of fifo systems to the Parikh coverability problem of serial asynchronous QDAS.

The proof of Theorem 4 relies on the next Lemma, that formalizes the relationship between reachable configurations of \mathcal{A}_F and reachable configurations of F .

For all $\gamma \in \Gamma$, we denote by s_γ^ℓ the location of \mathcal{TS}_γ that corresponds to line ℓ in Fig. 7. Then, we say that a configuration (G, \vec{d}) of \mathcal{A}_F encodes a configuration (s, w) of F , written $(G, \vec{d}) \triangleright (s, w)$ iff: (i) $s = \vec{d}(\text{state})$, (ii) G is of either shapes in Fig. 7 with $w = m_0 m_1 \dots m_n$, (iii) $\text{Parikh}(G)(s_{\text{main}}^8) = 1$ and (iv) there exists $m \in M \cup \{\varepsilon\}$ s.t. $\text{Parikh}(G)(s_m^{12}) = 1$. That is, s and w are encoded as described above, **main** is at line 8, and the running **m** block is at line 12. Then:

Lemma 5 *Let F be a FIFO system, let c be a configuration of F , and let \mathcal{A}_F be its associated QDAS. For all run $(s_0, w_0)(s_1, w_1) \dots (s_n, w_n)$ of F s.t. for all $0 \leq i < n$: $s_i \neq c$, there exists $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_F)$ s.t. $(G, \vec{d}) \triangleright (s_n, w_n)$. Moreover, for all $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_F)$ and for all configuration (s, w) of F : $(G, \vec{d}) \triangleright (s, w)$ implies $(s, w) \in \text{Reach}(F)$*

Proof. First, we consider a run $(s_0, w_0)(s_1, w_1) \dots (s_n, w_n)$ of F s.t. for all

$0 \leq i < n$: $s_i \neq c$, and build a run $(G_0, \vec{d}_0)a_0(G_1, \vec{d}_1)a_1 \cdots a_{k-1}(G_k, \vec{d}_k)$ of \mathcal{A}_F s.t. $(G_k, \vec{d}_k) \triangleright (s_n, w_n)$, by induction on the length of F 's run.

Base case $n = 0$: Consider the run of \mathcal{A}_F that consists in executing lines 5, 6, 7 of **main** (which sets the **head** variable to ε), then dequeuing the ε block from the queue, then executing lines 10 and 11 of ε . Remark that the test at line 10 is not satisfied, as **head** = ε , and that the queue is now empty. Clearly, the resulting configuration $(G, \vec{d}) \triangleright (s_F^0, w_0)$ as $w_0 = \varepsilon$.

Inductive case $n = \ell$. Let us assume that there is a reachable configuration (G, \vec{d}) of \mathcal{A}_F s.t. $(G, \vec{d}) \triangleright (s_{\ell-1}, w_{\ell-1})$, and let us build a sequence of \mathcal{A}_F transitions that is fireable from (G, \vec{d}) and reaches a configuration encoding (s_ℓ, w_ℓ) . In (G, \vec{d}) , there is, by definition of \triangleright , a task running a b block, for $b \in M \cup \{\varepsilon\}$, that is at line 12. Moreover, $\vec{d}(\text{state}) = s_{\ell-1}$. Let δ be the transition of F s.t. $(s_{\ell-1}, w_{\ell-1}) \xrightarrow{\delta} (s_\ell, w_\ell)$. By hypothesis, $s_{\ell-1} \neq c$, hence, we let b execute line 12; select $\delta = (s_{\ell-1}, a, s_\ell)$ at line 13; execute line 14, where the condition of the **if** is not satisfied as $s = s_{\ell-1} = \text{state}$; and execute line 16, which reaches a configuration (G', \vec{d}') where $\vec{d}'(\text{state}) = s_\ell$. We consider three cases to complete the simulation of δ in \mathcal{A}_F . If $a = !n$, the b task performs an asynchronous dispatch of n to \mathbf{q} , and jumps to line 11, then 12. Clearly, the resulting configuration (G'', \vec{d}'') is s.t. $(G'', \vec{d}'') \triangleright (s_\ell, w_\ell)$ (in particular, the dispatch has correctly updated the content of the queue). If $a = \varepsilon$, the b task jumps directly to line 11, then to line 12. Again, the resulting configuration (G'', \vec{d}'') is s.t. $(G'', \vec{d}'') \triangleright (s_\ell, w_\ell)$, as the content of the queue has not been modified. Finally, if $a = !n$, the running b block sets **head** to n and terminates. Let (G'', \vec{d}'') be the \mathcal{A}_F configuration reached at that point. As δ is fireable from $(s_{\ell-1}, w_{\ell-1})$ in F , since $(G, \vec{d}) \triangleright (s_{\ell-1}, w_{\ell-1})$, and as the content of the queue has not been modified since then, the head of \mathbf{q} is necessarily an n block in G'' . Moreover, $\vec{d}''(\text{head}) = n$ and $\vec{d}''(\text{state}) = s_\ell$. Thus, we let the scheduler dequeue this n block, and we let the task running it execute line 10 (where the condition of the **if** is not satisfied), then line 11. Clearly, the resulting configuration encodes (s_ℓ, w_ℓ) .

Now, let $\rho = (G_0, \vec{d}_0)a_0(G_1, \vec{d}_1)a_1 \cdots a_{n-1}(G_n, \vec{d}_n)$ be a run of \mathcal{A}_F s.t. there is (s, w) with $(G_n, \vec{d}_n) \triangleright (s, w)$, and let us build, by induction on the length of this run, a run $(s_F^0, w_0)(s_1, w_1) \cdots (s_k, w_k)$ a run of F s.t. $(s_k, w_k) = (s, w)$.

Let $K = |\{(G_i, \vec{d}_i) \mid \text{Parikh}(G_i)(s_m^{12}) = 1 \text{ for } m \in M \cup \{\varepsilon\}\}|$, i.e., K is the number of times an m block reaches line 12 along ρ . Let us consider the increasing monotonic function $\rho : \{1, \dots, K\} \rightarrow \{0, \dots, n\}$ s.t. for all $0 \leq i \leq n$: there exists $m \in M \cup \{\varepsilon\}$ s.t. $\text{Parikh}(G_i)(s_m^{12}) = 1$ iff there is $1 \leq j \leq K$ s.t. $G_i = \rho(j)$, that is, $\rho(i)$ is the index, in ρ of the i th time a configuration is reached where an m block is at line 12. Clearly, by definition of \triangleright only the $(G_{\rho(j)}, \vec{d}_{\rho(j)})$ configurations (for $1 \leq j \leq K$) can encode a configuration of F , as no m block is at line 12 in the other configurations of ρ . So, it is sufficient to show that all those $(G_{\rho(j)}, \vec{d}_{\rho(j)})$ configurations encode a reachable configuration of F . We proceed by induction on j , and show that: for all $1 \leq j \leq K$: $(G_{\rho(j)}, \vec{d}_{\rho(j)})$

encodes a reachable configuration of F and $G_{\rho(j)}$ contains exactly one m task (for $m \in M \cup \{e\}$), that has been dequeued from \mathbf{q} .

Base case $j = 0$: Observe that the subrun $(G_0, \vec{d}_0) a_0 \cdots a_{\rho(1)-1} (G_{\rho(1)}, \vec{d}_{\rho(1)})$ is necessarily an initialization phase where **main** sets **state** to s_F^0 , **head** to ε , dispatches an ε block, and reaches line 8, where it will stay forever. Then, the scheduler dequeues the ε block, which empties the queue. The ε task then traverses line 10 (as **head** = ε) and 11 and reaches line 12. So, clearly $(G_{\rho(0)}, \vec{d}_{\rho(0)}) \triangleright (s_F^0, \varepsilon)$ and contains exactly one m task (for $m \in M \cup \{e\}$), that has been dequeued from \mathbf{q} .

Inductive case $j = \ell$: Let us assume that $(G_{\rho(\ell-1)}, \vec{d}_{\rho(\ell-1)})$ encodes a reachable configuration $(s_{\ell-1}, w_{\ell-1})$ of F . We consider several cases. If $(G_{\rho(\ell-1)}, \vec{d}_{\rho(\ell-1)}) = (G_{\rho(\ell)}, \vec{d}_{\rho(\ell)})$ we are done. Otherwise, we have necessarily performed one iteration (possibly interrupted at line 12, 14 or 19) of the while loop at line 11 between $(G_{\rho(\ell-1)}, \vec{d}_{\rho(\ell-1)})$ and $(G_{\rho(\ell)}, \vec{d}_{\rho(\ell)})$, as, by induction hypothesis, $G_{\rho(\ell-1)}$ contains exactly one m task (with $m \in M \cup \{\varepsilon\}$) that blocks \mathbf{q} , and **main** can only loop at line 8, which does not modify the current configuration. Then, observe that the conditions of the **if** at lines 12 and 14 were necessarily false during the iteration. Otherwise, m would have reached line 21, from which it cannot escape. From that point, no configuration is reachable where an **m** block is at line 12, and $(G_{\rho(\ell)}, \vec{d}_{\rho(\ell)})$ cannot exist. Thus, we consider three cases:

- If we have entered the **if** at line 16 during the iteration, then a transition of the form $(s, !n, s')$ has been guessed, with **state** = s and a dispatch of n has been performed into q . As $(G_{\rho(\ell-1)}, \vec{d}_{\rho(\ell-1)}) \triangleright (s_{\ell-1}, w_{\ell-1})$ by induction hypothesis, $s_{\ell-1} = s$, and thus $(s, !n, s')$ is fireable from $(s_{\ell-1}, w_{\ell-1})$ and reaches $(s', n \cdot w_{\ell-1})$. Clearly, this configuration is encoded by $(G_{\rho(\ell)}, \vec{d}_{\rho(\ell)})$.
- If we have entered the **else if** at line 17 during the iteration, then a transition of the form $(s, ?n, s')$ has been guessed, with **state** = s , **head** has been set to n , the current m block has been terminated, a new block m' has been dequeued by the scheduler (as there is necessarily a running **m** block in $G_{\rho(\ell)}$). Moreover $m' = n$, because m' has to be at line 12 in $G_{\rho(\ell)}$, so the test of line 10 had to be false to allow m' to reach line 12. As $(G_{\rho(\ell-1)}, \vec{d}_{\rho(\ell-1)}) \triangleright (s_{\ell-1}, w_{\ell-1})$ by induction hypothesis, $s_{\ell-1} = s$. As a dequeue of a block $m' = n$ has been performed, $w_{\ell-1}$ is of the form $w \cdot n$. Thus, $(s, ?n, s')$ is fireable from $(s_{\ell-1}, w_{\ell-1})$ and reaches (s', w) . Clearly, this configuration is encoded by $(G_{\rho(\ell)}, \vec{d}_{\rho(\ell)})$.
- Finally, if neither the **if** nor the **else if** have been entered during the iteration, then a transition of the form (s, ε, s') has been guessed, with **state** = s . As $(G_{\rho(\ell-1)}, \vec{d}_{\rho(\ell-1)}) \triangleright (s_{\ell-1}, w_{\ell-1})$ by induction hypothesis, $s_{\ell-1} = s$, and thus (s, ε, s') is fireable from $(s_{\ell-1}, w_{\ell-1})$ and reaches $(s', w_{\ell-1})$. Clearly, this configuration is encoded by $(G_{\rho(\ell)}, \vec{d}_{\rho(\ell)})$. \square

We can now prove Theorem 4: *Proof.* Let F be a FIFO system, with set of messages M and associated serial asynchronous QDAS \mathcal{A}_F and let c be a control location of F . For all $m \in M \cup \{\varepsilon\}$, let f_m be the Parikh image s.t. $f_m(s_{\text{main}}^{21}) = 1$ and $f_m(s) = 0$ for all $s \neq s_{\text{main}}^{21}$. Remark that there are only finitely many such f_m . Then, we show that c is reachable in F iff there exists $m \in M \cup \{\varepsilon\}$ s.t. f_m is Parikh-coverable in \mathcal{A}_F .

Assume c is reachable in F , and let (c, w) be a configuration in $\text{Reach}(F)$. Without loss of generality, assume c is reachable by run that visits c only once. By Lemma 5, there is $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_F)$ s.t. $(G, \vec{d}) \triangleright (c, w)$. Hence, in (G, \vec{d}) , there is a task running an m block (for $m \in M \cup \{\varepsilon\}$) that is at line 12, and $\vec{d}(\text{state}) = c$. Thus, m can execute one step and reach line 21, so f_m is Parikh coverable in \mathcal{A}_F .


For the reverse direction, assume there is $m \in M \cup \{\varepsilon\}$ that is Parikh-coverable in \mathcal{A}_F . Hence, there is $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_F)$ where a task running block m is at line 21. The only way for that block to reach line 21 is from line 12, with a valuation \vec{d}' s.t. $\vec{d}'(\text{state}) = c$. Thus, there is, in $\text{Reach}(\mathcal{A}_F)$ a configuration (G', \vec{d}') with $\vec{d}'(\text{state}) = c$, a task running an m block at line 12, and necessarily `main` at line 8 (otherwise, only `main` would be running). Hence, (G', \vec{d}') is a reachable configuration of \mathcal{A}_F s.t. $(G', \vec{d}') \triangleright (c, w)$ for some queue content w . Thus, by Lemma 5, $(c, w) \in \text{Reach}(F)$, and c is reachable in F .

We have thus reduced the control location reachability problem of FIFO systems to the Parikh coverability problem of serial asynchronous QDAS (using only one serial queue). The former is undecidable. Hence the theorem. \square

B.3 Concurrent Qdas

We reduce the reachability problem of two counter systems. Let us give the intuition of the construction. For each \mathcal{P} , we construct a QDAS $\mathcal{A}_{\mathcal{P}}$ s.t. all reachable CTG in $\mathcal{A}_{\mathcal{P}}$ encode configurations of \mathcal{P} and are of the form depicted in Fig. 8. That is, (after an initialization phase), there are always three tasks that are unblocked: a *main* task to simulate \mathcal{P} 's control structure, and, for each $i = \{1, 2\}$, either a task *eins*(i) or a task *null*(i). If the task *null*(i) is unblocked, then counter i is zero in the current configuration of \mathcal{P} . Otherwise, the current valuation of counter i is encoded by the number of *eins*(i) tasks in the CTG. Remark that, as in the case of synchronous QDAS, the parts of the CTG that encode each counter behave as pushdown stacks. Finally, the control location of \mathcal{P} is recorded in global variable `state`.

The actual operations on the counters will be simulated by the *eins*(i) and *null*(i) running tasks. As *main* simulates the control structure, we need to synchronize *main* with those *eins*(i) and *null*(i) tasks. Let us explain intuitively how we can achieve *rendezvous* synchronization between running tasks using global variables of QDAS. Consider a QDAS with three global variables ℓ_1, ℓ_2 ranging over Boolean and X over a finite set of ‘messages’ M . Let γ_1 and γ_2 be two blocks whose LTS are:



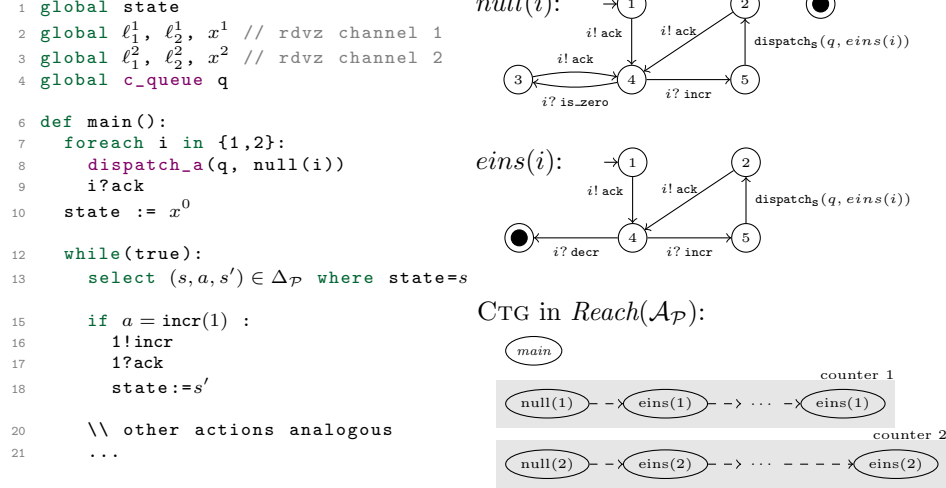



Figure 8: From a two counter system to a QDAS: *main* and *null(i)*, *eins(i)* for $i = 1, 2$



Assume a configuration c of the QDAS where $\ell_1 = \ell_2 = 0$ and where two distinct tasks are running γ_1 and γ_2 , are unblocked, and are in s_0 and s'_0 respectively. Assume that no other task can access ℓ_1, ℓ_2 and m . It is easy to check that, from c , there is only one possible interleaving of the transitions of γ_1 and γ_2 . So if γ_2 reaches s'_5 from c , then γ_1 must have reached s_5 , and the $x = m$ test in γ_1 has been fired *after* the $x \leftarrow m$ assignment in γ_2 . This achieves rendezvous synchronisation between γ_1 and γ_2 , with the passing of message m . This can easily be extended to rendezvous via different “channels”, by adding extra global variables. So, we extend the syntax of QDAS by allowing transitions of the form $(s_0, c!m, s_5)$ and $(s'_0, c?m, s'_5)$ (for $m \in M$) to denote respectively a send and a receive of message m on a rendezvous channel c .

We rely on this mechanism to let *main* send operations to be performed on the counters to the *null(i)* and *eins(i)* running tasks. More precisely, for a 2Cs $\mathcal{P} = \langle X, x^0, \Sigma_P, \Delta_P \rangle$, we build the QDAS $\mathcal{A}_P = \langle CQID, \emptyset, \Gamma, \text{main}, \mathcal{X}, \Sigma, (\mathcal{T}\mathcal{S}_\gamma)_{\gamma \in \Gamma} \rangle$ where $CQID = \{q\}$, $\Gamma = (\{null, eins\} \times \{1, 2\}) \cup \{main\}$, $\mathcal{X} = \{\ell_1^1, \ell_2^1, x^1, \ell_1^2, \ell_2^2, x^2\}$ where x^1, x^2 range over the domain $\{incr, decr, is_zero, ack\}$, and the transition systems are given in Fig.8. The variables \mathcal{X} encode two channels that we call 1 and 2 in the pseudo code of Fig. 8. The *main* task runs an infinite `while` loop (line 12 onwards) that consists in guessing a transition (s, a, s') of F and synchronising, via *rendezvous* on the channels 1 and 2, with the relevant *null* or *eins* unblocked task, to let it execute the operation on the counter. When a *null(i)* or *eins(i)* receives an `incr` message, it performs an asynchronous dispatch of *eins(i)* into q to increment counter i , and acknowledges the operation to *main*, thanks to message `ack`. When an *eins* block receives a `decr` message, it terminates, which decrements the counter. *null* blocks cannot


receive **decr** messages, so, if *main* requests a **decr** operation when the counter is zero, *main* gets blocked. This means that the guessed transition was not fireable in the currently simulated 2Cs configuration, and ends the simulation. Finally, only *null* blocks can receive and acknowledge **is_zero** messages, so, again, *main* is blocked after sending **is_zero** to a non-zero counter. Note that we need both *asynchronous* calls to start two counters in parallel, and *synchronous* calls to encode the counter values. The result of Theorem 5 follows directly from:

Proposition 10 *Given a 2Cs, then we can reduce its reachability question to the Parikh coverability question for a concurrent QDAS that demands both synchronous and asynchronous dispatch actions.*

As discussed before, we can separate each G for $(G, \vec{d}) \in \text{Reach}(\mathcal{A}_C)$ into three components, one consisting only of a vertex v_0 with $\lambda(v_0) = \text{main}$ and two paths $v_1 v_2 \dots v_k$ and $v'_1 v'_2 \dots v'_l$ which we will call *counter1* and *counter2* in the following.

As before, we define a relation between configurations of the 2Cs \mathcal{C} and the QDAS \mathcal{A}_{C_c} . For $c = (G, \vec{d}) \in \text{Reach}(\mathcal{A}_C)$ and $y = (x, k, l) \in \text{Reach}(\mathcal{C}) \subseteq X \times \mathbb{N} \times \mathbb{N}$ we write $c \triangleright y$ if $\vec{d}(\text{state}) = x$, $|\text{counter1}| = k$, and $|\text{counter2}| = l$.


The rendezvous assures a unique interleaving of actions of *main*, *null(1)*, and *null(2)* until *main* reaches line 12. Let us in the following consider the reached configuration $c^0 = (G^0, \vec{d}^0)$ with $\vec{d}^0(\text{state}) = x^0$, $\vec{d}^0(\ell_0) = \vec{d}^0(\ell_1) = 0$ and G^0 with




(where s_{12} is the state of *main* in line 12) as “initial” configuration of the QDAS.

Note that *counter1* and *counter2* are independent, i.e., they do not synchronize except via *main*. Further, there is no more than one task active in *counter1* and *counter2*. The unique tasks *zero(1)* and *zero(2)* never terminate. The rendezvous synchronization assures that there is only *one* possible interleaving between the *main* task and the currently running tasks in *counter1* and *counter2*:

- *main* does loops of the form



- which leads to the following interleaving of actions of *main* with actions of the i -th counter component.



where $\bigcirc \xrightarrow{c_i(a)} \bigcirc$ translates the sent action a to a meta-action $c_i(a)$ of the i -th counter as follows:

- an action **incr** is mapped to the action $\text{dispatch}_s(q, \text{eins}(i))$ and the activation of the dispatched task

- an action **decr** is mapped to the termination of the current task; which is only possible if the current task is a block $ains(i)$
- the test for empty stack is mapped to an epsilon action; this action is only possible in $null(i)$.

Note that if $c_i(a)$ is not possible, then there will be no acknowledgement, hence \mathcal{A}_C blocks.

Thus we can cut a run of \mathcal{A}_C into (an initial phase and) a sequence of phases of the above form that will be abbreviated $trans(s, a, s')$ in the following.

Lemma 6 *Let \mathcal{C} be a 2CS and \mathcal{A}_C the associated QDAS, if $y \in Reach(\mathcal{C})$ then there exists $c \in Reach(\mathcal{A}_C)$ such that $c \triangleright y$. Further, if $c = (G, \vec{d}) \in Reach(\mathcal{A}_C)$ where \vec{d} evaluates $\vec{d}(\ell_0) = \vec{d}(\ell_1) = 0$, then there exists $y \in Reach(\mathcal{C})$ with $c \triangleright y$.*

Proof.

Given a run $x_0 \delta_1 x_1 \delta_2 \dots \delta_k x_k$ of \mathcal{C} , then there exists a run of \mathcal{A}_C that can be cut into phases t_1, \dots, t_k where $t_i = trans(s_{i-1}, a_i, s_i)$ where $\delta_i = (s_{i-1}, a_i, s_i)$ for $1 \leq i \leq k$. Obviously $c^0 \triangleright x^0$ and $\vec{d}^0(\ell_0) = \vec{d}^0(\ell_1) = 0$. Hence, the reverse direction follows by a straightforward inductive argument.